

Frequency-Referenced Multiplexing for Energy-Efficient Data Center Optical Interconnects

Shea Keogh

Department of Computer Science, George Mason University, Fairfax, VA, USA.
keogh1992@gmu.edu

Mitchel Leddin

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

Scot Hanbury

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.
scot.hanbury@uc.edu

Tage Aberg

Department of Computer Science, University of New Hampshire, Durham, NH, USA.
tage1995@unh.edu

Abstract

Data center optical interconnects are increasingly constrained by the simultaneous demands of bandwidth density, latency determinism, energy proportionality, thermal feasibility, and multi-tenant fairness. Frequency-referenced multiplexing has emerged as a system-level strategy for coordinating many optical channels through a shared spectral and timing reference, potentially reducing tuning overhead, stabilizing channel spacing, and enabling guaranteed bandwidth with low latency. This paper examines frequency-referenced multiplexing not merely as a physical-layer technique but as an architectural and governance mechanism for energy-efficient data center fabrics. It analyzes how referenced multiplexing interacts with co-packaged optics, disaggregated compute, optical circuit switching, and service-level guarantees. The discussion emphasizes structural trade-offs among transceiver complexity, control-plane overhead, cooling constraints, failure containment, and infrastructure renewal. It also considers deployment pathways in brownfield and greenfield facilities, the implications for fair resource allocation across tenants, and the policy questions raised by energy reporting, carbon accounting, and public digital infrastructure. The central argument is that frequency-referenced multiplexing can improve energy efficiency only when its benefits are embedded in robust scheduling, interoperable standards, and accountable operational practices. Without such socio-technical alignment, gains at the link level may be absorbed by overprovisioning, thermal penalties, or opaque governance.

Keywords

data center optical interconnects; frequency-referenced multiplexing; energy efficiency; co-packaged optics; latency determinism; optical switching; sustainability; infrastructure governance.

1. Introduction

Data centers have become the computational substrate of contemporary economies, mediating cloud services, artificial intelligence training, financial clearing, scientific simulation, public administration, and consumer platforms. Their internal architecture is therefore no longer a narrow engineering concern but a matter of infrastructure policy, energy governance, and social dependency [1]. As computational demand continues to rise, the movement of data within and between racks, nodes, and accelerators has become a principal bottleneck. Electrical interconnects, once adequate for short-reach communication, now encounter severe limits in bandwidth density, signal equalization, and energy per bit as link rates increase [2]. The slowdown of traditional device scaling further intensifies pressure on system architects to extract performance through parallelism, specialization, and interconnection rather than through isolated transistor improvements [3]. Optical interconnects have consequently moved from an exotic alternative to a central design consideration in warehouse-scale machines and high-performance computing systems [4], [5].

Frequency-referenced multiplexing offers a promising response to this pressure by coordinating multiple optical channels through a common frequency reference. Rather than treating each wavelength or channel as an independently tuned resource, the approach stabilizes and aligns channels relative to a shared reference, which can support guaranteed bandwidth and low latency in multiplexed communication [6]. The significance of this capability extends beyond link-level performance. It touches scheduling, isolation, thermal management, failure recovery, and the economics of deploying optical infrastructure at scale. If channels can be coordinated with greater precision and less per-channel overhead, operators may achieve more predictable service levels while reducing the energy spent on tuning, retuning, and overprovisioning. Yet the same capability also creates new dependencies: on reference distribution, on control-plane coordination, on standards, and on the institutional capacity to manage shared spectral resources fairly.

This paper argues that frequency-referenced multiplexing should be evaluated as a socio-technical architecture rather than only as a photonic technique. The central question is not whether referenced multiplexing can deliver higher bandwidth, but under what structural conditions it can deliver energy efficiency without sacrificing robustness, fairness, or sustainability. Section 2 examines the architectural pressures that make optical multiplexing necessary in modern data centers. Section 3 conceptualizes frequency-referenced multiplexing as a coordination primitive. Section 4 analyzes energy and thermal trade-offs. Section 5 considers robustness, fairness, and governance. Section 6 explores deployment pathways and infrastructure transitions. Section 7 discusses sustainability and policy implications. The conclusion synthesizes the argument and identifies conditions for responsible adoption.

2. System Context and Architectural Pressures

Modern data centers are best understood as systems of systems in which compute, storage, memory, and network resources are distributed across a hierarchy of racks, rows, buildings, and geographic regions. The dominant traffic pattern is no longer the north-south flow between users and servers but the east-west flow among servers, storage nodes, and accelerators [8]. This shift has made the data center network a primary determinant of application performance, energy consumption, and operational cost [9]. Scale-out architectures rely on many commodity or semi-custom nodes connected through high-radix switches, yet the electrical signaling required to sustain such fabrics becomes increasingly power-hungry and difficult to cool as line rates rise [10]. The result is a structural tension: the demand for more bandwidth collides with the thermal and power limits of the facility.

Optical interconnects mitigate this tension by carrying signals over fiber with lower loss and greater reach than copper, while also offering wavelength-division multiplexing as a path to higher aggregate capacity. However, optical links are not automatically energy-efficient. Lasers, modulators, drivers, receivers, and thermal control circuits all consume power, and the electro-optical conversions at each end introduce overhead [11]. The challenge is therefore to design optical fabrics that reduce conversions, simplify tuning, and match capacity to actual demand. Co-packaged optics, silicon photonics, and advanced modulators have advanced this agenda, but they also increase integration complexity and create new failure modes [12], [13], [14], [15]. Frequency-referenced multiplexing enters this context as a method for making multiplexed optical resources more deterministic and controllable.

The broader system context also includes the rise of accelerators and specialized hardware. Reconfigurable fabrics and domain-specific accelerators can deliver substantial performance gains, but they intensify data movement because specialized units must be fed and coordinated [18]. Meanwhile, power usage effectiveness and carbon reporting have become governance concerns, not only engineering metrics [7]. As operators face scrutiny over energy and water use, the efficiency of internal interconnects becomes part of public accountability. Frequency-referenced multiplexing therefore must be assessed in relation to both technical performance and institutional legitimacy. A link-level improvement that increases total facility energy or undermines fair access would be a poor architectural choice.

3. Frequency-Referenced Multiplexing as a Coordination Primitive

Frequency-referenced multiplexing can be understood as a coordination primitive because it organizes multiple optical channels around a shared spectral and temporal reference. In conventional wavelength-division multiplexing, each channel may be independently generated, tuned, and stabilized, which creates cumulative complexity as channel counts grow. A common reference can reduce drift, simplify channel assignment, and support tighter guarantees on bandwidth and latency [6]. The system-level implication is significant: optical resources become less like a collection of independent links and more like a managed pool whose properties are known to the scheduler and control plane. This can improve predictability for latency-sensitive applications and enable finer-grained resource isolation.

The approach also resonates with broader trends in optical networking, where control planes increasingly manage spectrum, topology, and restoration across multiple layers [19]. If frequency referencing provides a stable basis for channel coordination, then higher-level schedulers can allocate capacity with greater confidence. For example, a tenant requiring deterministic latency for distributed training or financial workloads could be assigned referenced channels that are reserved and monitored as a service class. Other tenants might receive best-effort or elastic capacity. This creates opportunities for fairness and service differentiation, but it also requires governance: someone must define the reference, allocate spectrum, resolve conflicts, and audit compliance. The technical primitive does not by itself produce equitable outcomes.

Frequency-referenced multiplexing must also be distinguished from simple wavelength parallelism. The key is not merely that multiple frequencies exist, but that they are referenced to a common standard that enables coordinated behavior. This can reduce the need for per-channel tuning loops and may allow transceivers to operate with lower overhead. At the same time, it introduces a single point of coordination that must be robust. If the reference distribution fails, many channels may be affected simultaneously. Architectural resilience therefore requires redundant references, graceful degradation, and fallback modes in which

channels can operate independently or with reduced guarantees. The trade-off between coordination efficiency and systemic fragility is central to deployment decisions.

4. Energy Efficiency and Thermal Management

The energy efficiency case for frequency-referenced multiplexing rests on reducing the work required to establish, maintain, and recover optical channels. In large data centers, a significant share of interconnect power is consumed not only by transmission but by tuning, retuning, signal processing, and thermal stabilization. If a shared reference reduces wavelength drift and simplifies channel locking, transceivers may spend less energy on control and compensation. This can improve energy proportionality, meaning that power consumption scales more closely with useful throughput. However, the net benefit depends on the reference generation and distribution overhead, the complexity of the transceiver design, and the extent to which referenced channels are actually utilized.

Silicon photonics and integrated modulators have made optical transceivers more compact and potentially cheaper, but they also exhibit temperature sensitivity and manufacturing variability [13], [14], [15]. Frequency referencing can help manage variability by giving many channels a common anchor, but it cannot eliminate thermal problems. Lasers remain sensitive to temperature, and dense co-packaged optics can concentrate heat near high-power switching and compute components. Thermal management therefore becomes a co-design problem: the optical architecture, packaging, cooling system, and workload scheduler must be optimized together. A referenced multiplexing scheme that improves spectral efficiency but increases cooling load may not improve total facility energy performance [16].

Energy-aware traffic engineering offers a complementary layer of optimization [20]. If the network control plane knows which channels are referenced and how much energy they consume under different loads, it can route traffic to minimize energy while respecting latency and bandwidth guarantees. This requires measurement and telemetry at a granularity that many data centers currently lack. It also raises questions about rebound effects: efficiency gains may encourage greater consumption, leading to higher absolute energy use. Without governance mechanisms such as carbon-aware scheduling, efficiency improvements can be absorbed by growth. Frequency-referenced multiplexing should therefore be deployed with transparent energy accounting and workload policies that treat efficiency as a means to sustainability rather than as an end in itself.

5. Robustness, Fairness, and Governance

Robustness is a first-order concern for any shared coordination mechanism. Frequency-referenced multiplexing depends on a reference that many channels trust. If that reference is disrupted, the effects may propagate across tenants and services. A robust design must therefore include redundancy, anomaly detection, isolation, and recovery procedures. It must also specify how guarantees degrade under failure. For example, if a referenced channel loses synchronization, does it fall back to a lower-rate independent mode, or does it fail closed to protect other channels? These choices are not purely technical; they allocate risk among operators, tenants, and end users. Governance frameworks should make such allocations visible and contestable.

Fairness is equally important. Data center operators increasingly serve multiple tenants with different willingness to pay, latency requirements, and political sensitivity. Frequency-referenced multiplexing could support fair allocation by making guaranteed bandwidth and low latency explicit services [6]. But without oversight, it could also entrench priority for the

highest-paying customers, leaving public-interest workloads with degraded service. Fairness requires scheduling policies that consider not only price but also service criticality, legal obligations, and historical inequities in digital access. The technical ability to partition spectrum does not determine the normative criteria for partition. Those criteria must be established through policy, contract, and operational governance.

Governance also extends to standards and interoperability. If frequency referencing is implemented through proprietary interfaces, operators may face lock-in, and smaller providers may be excluded from the market. Open standards can reduce barriers and enable multi-vendor ecosystems, but they also slow innovation and require compromise. The history of optical networking suggests that successful technologies often depend on layered standards that separate physical references, control protocols, and service definitions [19]. Data center operators, equipment vendors, standards bodies, and regulators all have roles to play. The central challenge is to preserve the efficiency benefits of coordination while preventing the concentration of control that coordination can enable.

6. Deployment Pathways and Infrastructure Transitions

Deployment of frequency-referenced multiplexing will occur along multiple pathways. Greenfield data centers can be designed from the outset with co-packaged optics, integrated references, and software-defined control planes. Brownfield facilities, by contrast, must retrofit optical systems into existing power, cooling, and cabling environments. The economics of retrofit are often decisive: if the new technology requires extensive replacement of racks, fibers, or cooling infrastructure, adoption may be delayed even when link-level benefits are clear [4], [5]. Modular deployment, in which referenced multiplexing is introduced first in high-value interconnects such as accelerator clusters or spine links, may offer a pragmatic transition path.

Disaggregation and composability further complicate deployment. In disaggregated architectures, compute, memory, and storage are pooled and connected through high-bandwidth optical fabrics. Frequency referencing could help manage the many channels required to connect pooled resources, but it also increases dependence on the optical control plane. If the control plane fails, resource pooling may collapse into isolated islands. Deployment planning must therefore include failure domains, fallback topologies, and service-level objectives that reflect the new dependencies. The same is true for reconfigurable accelerators, which may require rapid changes in interconnect capacity as workloads shift [18].

Infrastructure transitions also have labor and organizational dimensions. Operating a referenced optical fabric requires skills in photonics, network control, thermal management, and security. Many organizations currently separate these functions, which can lead to fragmented responsibility. Successful deployment may require new roles, cross-training, and integrated operations centers. Procurement practices must also evolve to evaluate total cost of ownership, energy use, and interoperability rather than focusing narrowly on peak bandwidth. Without such organizational adaptation, technical capabilities may remain underused or poorly governed.

7. Sustainability and Policy Implications

The sustainability implications of frequency-referenced multiplexing are both promising and ambiguous. On one hand, more efficient optical interconnects can reduce energy per bit and enable higher utilization of compute resources, potentially lowering total energy consumption.

On the other hand, efficiency gains can stimulate demand, leading to larger data centers and higher absolute energy use [7]. Policy must therefore focus on absolute emissions, water consumption, and embodied carbon, not only on operational efficiency. Frequency referencing may reduce the need for some retuning and thermal control, but it may also require additional reference generation and monitoring equipment. Life-cycle assessment is necessary to determine whether the net environmental effect is positive.

Policy frameworks should also address transparency. Operators should be required or encouraged to report interconnect energy consumption, utilization, and service-level performance in ways that allow comparison and accountability. Carbon-aware scheduling can shift workloads to times and places where low-carbon energy is available, but this requires accurate telemetry and flexible workloads [20]. Frequency-referenced multiplexing can support such scheduling by providing deterministic channels with known energy characteristics. However, the benefits will be limited if reporting remains voluntary or inconsistent. Regulators, standards bodies, and industry consortia all have roles in establishing common metrics.

Finally, sustainability is intertwined with fairness. Communities that host data centers often bear environmental burdens such as water use, land consumption, and grid strain, while the benefits of digital services may be distributed unevenly. Energy-efficient interconnects can reduce some local impacts, but they do not resolve questions of siting, taxation, or community consent. A just transition requires that efficiency improvements be paired with participatory governance and public accountability. Frequency-referenced multiplexing is therefore not only a question of photons and packaging; it is a question of who benefits, who bears risk, and how infrastructure decisions are made.

8. Conclusion

Frequency-referenced multiplexing offers a credible path toward more energy-efficient and deterministic data center optical interconnects. By coordinating multiple channels through a shared reference, it can reduce tuning overhead, improve spectral stability, and support guaranteed bandwidth with low latency [6]. These capabilities align with the architectural needs of modern data centers, where east-west traffic, accelerators, disaggregation, and multi-tenancy place intense pressure on interconnect capacity and predictability. The technology is most promising when integrated with co-packaged optics, silicon photonics, and software-defined control planes that can exploit its coordination benefits.

Yet the paper has argued that link-level efficiency is not sufficient. The structural trade-offs among reference dependency, control-plane complexity, thermal load, failure containment, and interoperability shape whether frequency-referenced multiplexing improves total system performance. Fairness and governance are equally critical. Guaranteed bandwidth and low latency can reinforce equitable service classes, but they can also entrench privilege if allocation criteria are opaque or market-driven alone. Robustness requires redundancy and graceful degradation, while sustainability requires absolute energy accounting and life-cycle thinking.

The policy implication is that frequency-referenced multiplexing should be developed within an open, accountable, and interdisciplinary framework. Standards should support interoperability; operators should report energy and performance transparently; regulators should consider absolute environmental impacts; and researchers should study both technical and social dimensions of deployment. Under these conditions, frequency-referenced

multiplexing can contribute to data center infrastructures that are not only faster and more efficient but also more resilient, fair, and sustainable.

References

1. Barroso, L. A., Clidaras, J., & Hölzle, U. (2018). *The datacenter as a computer: Designing warehouse-scale machines* (3rd ed.). Morgan & Claypool.
2. Miller, D. A. B. (2009). Device requirements for optical interconnects to silicon chips. *Proceedings of the IEEE*, 97(7), 1166–1185.
3. Shalf, J. (2020). The future of computing beyond Moore's Law. *Philosophical Transactions of the Royal Society A*, 378(2166), 20190061.
4. Benner, A. F., Ignatowski, M., Kash, J. A., Kuchta, D. M., & Ritter, M. B. (2005). Exploitation of optical interconnects in future server architectures. *IBM Journal of Research and Development*, 49(4/5), 755–775.
5. Kachris, C., & Tomkos, I. (2012). A survey on optical interconnects for data centers. *IEEE Communications Surveys & Tutorials*, 14(4), 1021–1036.
6. Zhou, Z., Wei, J., Luo, Y., Clark, K. A., Sillekens, E., Deakin, C., ... & Liu, Z. (2023). Communications with guaranteed bandwidth and low latency using frequency-referenced multiplexing. *Nature Electronics*, 6(9), 694–702.
7. Masanet, E., Shehabi, A., Lei, N., Smith, S., & Koomey, J. (2020). Recalibrating global data center energy-use estimates. *Science*, 367(6481), 984–986.
8. Al-Fares, M., Loukissas, A., & Vahdat, A. (2008). A scalable, commodity data center network architecture. *ACM SIGCOMM Computer Communication Review*, 38(4), 63–74.
9. 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.
10. Vahdat, A., Al-Fares, M., Farrington, N., Mysore, R. N., Porter, G., & Radhakrishnan, S. (2010). Scale-out networking in the data center. *IEEE Micro*, 30(4), 29–41.
11. 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.
12. Sun, C., Wade, M. T., Lee, Y., Orcutt, J. S., Alloatti, L., Georgas, M. S., ... & Stojanović, V. M. (2015). Single-chip microprocessor that communicates directly using light. *Nature*, 528(7583), 534–538.
13. Reed, G. T., Mashanovich, G., Gardes, F. Y., & Thomson, D. J. (2010). Silicon optical modulators. *Nature Photonics*, 4(8), 518–526.
14. Liu, K., Ye, C. R., Khan, S., & Sorger, V. J. (2015). Review and perspective on ultrafast wavelength-size electro-optic modulators. *Laser & Photonics Reviews*, 9(2), 172–194.
15. Assefa, S., Shank, S., Green, W., Khater, M., Kiewra, E., Reinholm, C., ... & Vlasov, Y. (2012). A 90nm CMOS integrated nano-photonics technology for 25Gbps WDM optical communications applications. In *2012 International Electron Devices Meeting* (pp. 33.8.1–33.8.3). IEEE.

16. Hinton, K., Baliga, J., Feng, M., Ayre, R., & Tucker, R. S. (2011). Power consumption and energy efficiency in the internet. *IEEE Network*, 25(2), 6–12.
17. Hamilton, J. (2009). Cooperative expendable micro-slice servers (CEMS): Low cost, low power servers for internet-scale services. In *Proceedings of the 4th Biennial Conference on Innovative Data Systems Research (CIDR 2009)*.
18. Putnam, A., Caulfield, A. M., Chung, E. S., Chiou, D., Constantinides, K., Demme, J., ... & Burger, D. (2014). A reconfigurable fabric for accelerating large-scale datacenter services. *ACM SIGARCH Computer Architecture News*, 42(3), 13–24.
19. O'Mahony, M. J., Politi, C., Klonidis, D., Nejabati, R., & Simeonidou, D. (2006). Future optical networks. *Journal of Lightwave Technology*, 24(12), 4684–4696.
20. Christodoulopoulos, K., Kontodimas, K., & Varvarigos, E. (2015). Energy-aware traffic engineering in optical networks. *Optical Switching and Networking*, 17, 1–12.