

Intelligent Wavelength and Frequency Assignment in Elastic Optical Networks

Yatan C. Maharaj

Department of Electrical Engineering and Computer Science, University of Kansas, Lawrence, KS, USA.

ymaharaj@ku.edu

Abstract

Elastic optical networks have emerged as a foundational infrastructure for accommodating heterogeneous, high-capacity, and latency-sensitive digital services within a shared optical substrate. Unlike conventional wavelength-division multiplexing systems that operate on a fixed frequency grid, elastic optical networks allocate spectrum in finer granularities and permit variable-width channels, thereby improving spectral efficiency and enabling more flexible service provisioning. At the center of this flexibility lies the problem of intelligent wavelength and frequency assignment, which coordinates routing, modulation, spectrum allocation, regeneration, and restoration across physical and logical layers. This paper examines wavelength and frequency assignment as a system-level challenge rather than a narrow combinatorial optimization problem. It argues that intelligent assignment must balance spectral efficiency against fragmentation, resilience, operational complexity, energy consumption, fairness, and governance. The discussion considers architectural foundations, control-plane intelligence, machine learning for monitoring and failure management, deterministic service guarantees, infrastructure investment, and policy implications. Particular attention is given to the structural trade-offs that arise when operators seek to maximize utilization while preserving robustness and equitable access. The paper also explores how automation, telemetry, and cross-domain coordination shape deployment pathways in national and regional networks. It concludes that intelligent wavelength and frequency assignment is best understood as a socio-technical capability, one that requires not only improved algorithms but also transparent governance, interoperable standards, sustainable infrastructure planning, and institutional accountability.

Keywords

elastic optical networks; wavelength assignment; frequency assignment; routing and spectrum allocation; artificial intelligence; network governance; sustainability; resilience; fairness; optical infrastructure.

1. Introduction

The evolution of optical transport has moved from rigid, fixed-grid wavelength-division multiplexing toward elastic optical networking, in which spectrum is treated as a flexible and dynamically allocatable resource. This transition responds to the growing diversity of traffic patterns, including high-capacity data-center interconnect, mobile backhaul, cloud computing, scientific data transfer, and deterministic industrial communication. Early work on spectrum-efficient elastic optical path networks established the architectural promise of flexible spectrum allocation and distance-adaptive transmission [1]. Subsequent analyses framed elastic optical networking as a new dawn for the optical layer, capable of aligning capacity with service demand more closely than fixed-grid systems [2]. Within this architecture,

wavelength and frequency assignment becomes a central coordination function, determining which portions of the optical spectrum are assigned to which paths, at what width, and under what quality-of-transmission constraints.

The intellectual significance of this problem lies in its multi-layered character. At the physical layer, assignment must respect frequency slot granularity, guard bands, laser tuning ranges, filter passbands, amplifier gain profiles, and nonlinear impairments. At the network layer, it must accommodate routing decisions, traffic grooming, regeneration placement, protection paths, and restoration capacity. At the operational layer, it must support service-level agreements, maintenance windows, failure recovery, and capacity planning. At the governance layer, it intersects with spectrum policy, competition, universal service, and infrastructure resilience. Consequently, intelligent assignment cannot be reduced to a single objective such as minimizing blocking probability or maximizing spectral efficiency. It is a systems problem that requires balancing technical performance with institutional, economic, and social considerations.

This paper develops a long-form analysis of intelligent wavelength and frequency assignment in elastic optical networks. It begins by reviewing architectural foundations and then reframes the assignment problem as a system-level resource coordination challenge. It examines the role of artificial intelligence, control-plane automation, and deterministic multiplexing in improving assignment decisions. It then considers robustness, sustainability, fairness, governance, and deployment pathways. The central argument is that intelligence in optical spectrum assignment must be understood as a distributed capability embedded in infrastructures, institutions, and operational practices, not merely as a computational technique applied to a static optimization problem.

2. Architectural Foundations of Elastic Optical Networks

Elastic optical networks differ from fixed-grid systems by replacing rigid wavelength channels with a flexible frequency grid. In a fixed-grid architecture, channels are typically assigned to a standardized grid with fixed spacing, which can lead to either stranded capacity or inefficient use of spectrum when traffic demands do not match the grid exactly. Elastic optical networks allow channel widths to vary according to demand, transmission reach, and modulation format. This flexibility improves spectral efficiency, but it also introduces constraints related to spectrum contiguity, continuity, and fragmentation. The assignment process must therefore determine not only whether a path can be established but also whether the required spectrum can be placed in a contiguous block along the entire route [3].

The enabling technologies of elastic optical networks include bandwidth-variable transponders, flexible-grid reconfigurable optical add-drop multiplexers, sliceable bandwidth-variable transponders, and advanced modulation techniques. These components allow the optical layer to adapt to changing traffic conditions and to allocate spectrum in units finer than conventional wavelength channels. Research on elastic bandwidth allocation in flexible orthogonal frequency-division multiplexing networks demonstrated how spectrum can be partitioned and reassigned dynamically [4]. Surveys of orthogonal frequency-division multiplexing-based elastic core optical networking further clarified the architectural options and control requirements for flexible spectrum management [7]. Together, these developments create a physical and logical substrate in which frequency assignment is not merely a labeling exercise but a mechanism for aligning transmission physics with service requirements.

The control plane is equally important. Elastic optical networks require mechanisms for disseminating topology, spectrum availability, impairment information, and policy constraints. Generalized multiprotocol label switching and software-defined networking approaches provide frameworks for path computation and resource reservation, while centralized or hierarchical controllers can optimize spectrum assignment across domains. However, centralization introduces scalability and trust challenges, especially in multi-operator environments. Routing and spectrum allocation research has therefore explored distributed, centralized, and hybrid architectures. The choice of control architecture affects how quickly assignment decisions can be made, how much telemetry is available, and how easily policies can be enforced [8].

Architectural trade-offs are unavoidable. Finer spectrum granularity improves packing efficiency but increases the complexity of filtering, monitoring, and management. Distance-adaptive modulation increases capacity but requires accurate quality-of-transmission estimation. Flexible transponders reduce spare-part variety but may increase software complexity and interoperability risk. Protection and restoration improve resilience but consume spectrum that could otherwise carry revenue-generating traffic. These trade-offs mean that intelligent assignment must be evaluated against a broad set of operational outcomes, including deployment cost, energy use, service availability, and fairness across customers and regions.

3. The Assignment Problem as a System-Level Resource Coordination Challenge

Routing and spectrum assignment in elastic optical networks is often described as an extension of routing and wavelength assignment. The routing component selects a path, while the spectrum assignment component selects a contiguous frequency range along that path. In practice, the problem is more complex because modulation format, regeneration, transponder capabilities, and quality-of-transmission constraints are interdependent. A path that is feasible under one modulation format may be infeasible under another. A spectrum block that is available on each link may still be unusable because of filter cascading or nonlinear interference. Early formulations of routing and spectrum assignment highlighted these interdependencies and the difficulty of achieving optimal solutions in realistic networks [3]. Studies of spectrum-sliced elastic optical path networks further demonstrated the importance of assignment policies for blocking performance and resource utilization [6].

System-level coordination requires that assignment decisions be made with awareness of traffic demand, service priority, and network state. Some services require guaranteed bandwidth and low latency, while others tolerate variable rates and moderate delay. Some traffic is highly dynamic, such as data-center synchronization or mobile fronthaul, while other traffic is scheduled and predictable, such as bulk scientific transfer. Intelligent assignment should therefore differentiate among service classes rather than treating all demands as equivalent. Frequency-referenced multiplexing has been shown to support communications with guaranteed bandwidth and low latency, illustrating how physical-layer design can reinforce deterministic service objectives [12]. This does not make assignment a single-purpose mechanism, but it shows that spectrum placement can be aligned with performance guarantees when architecture and control are co-designed.

Fragmentation is one of the most persistent challenges. As channels are established and torn down, available spectrum becomes divided into noncontiguous blocks. A new demand may require a contiguous block larger than any single gap, even though total free spectrum is sufficient. Fragmentation increases blocking probability and may force unnecessary

regeneration or rerouting. Surveys of spectrum management techniques have examined assignment, defragmentation, and reconfiguration strategies for mitigating this problem [5]. However, defragmentation itself can cause service disruption, requiring careful scheduling and risk assessment. Intelligent assignment must therefore consider not only current requests but also future demand patterns and the cost of reconfiguration.

Fairness enters the assignment problem through the distribution of spectrum access. In a purely utilization-maximizing regime, large or high-priority customers may capture the best spectrum blocks, while smaller or remote users experience higher blocking. Over time, such outcomes can entrench inequalities in service quality and infrastructure investment. A system-level approach should therefore incorporate fairness constraints or policy weights into assignment decisions. This does not mean abandoning efficiency, but it does mean recognizing that optical spectrum is a shared infrastructure resource with public and commercial dimensions. The assignment problem is thus simultaneously technical, economic, and normative [18].

4. Intelligence, Automation, and Control-Plane Design

Intelligence in wavelength and frequency assignment can be understood across several time scales. At the shortest time scale, assignment must respond to instantaneous link conditions, failures, and traffic fluctuations. At the medium time scale, it must support capacity planning, maintenance, and reconfiguration. At the longest time scale, it must inform infrastructure investment, spectrum policy, and technology migration. Artificial intelligence and machine learning can contribute at each scale, but their value depends on data quality, model interpretability, and integration with control-plane mechanisms. Machine learning for failure management in optical networks has been explored as a way to predict degradations, localize faults, and trigger restoration before service impact occurs [10]. Broader reviews of artificial intelligence in optical networks have mapped applications ranging from traffic prediction to optical performance monitoring and autonomous operation [11].

The control plane provides the operational context in which intelligent assignment occurs. A centralized software-defined networking controller may collect telemetry from transponders, amplifiers, and monitors, then compute spectrum assignments using optimization or learning-based methods. A distributed control plane may allow faster local decisions but provide less global visibility. Hierarchical architectures attempt to combine both, with domain controllers managing local spectrum and a higher-level orchestrator coordinating inter-domain paths. The design choice affects robustness, scalability, and governance. If intelligence is centralized, operators must address single points of failure and data sovereignty. If intelligence is distributed, they must ensure consistency and avoid conflicting assignments [8].

Machine learning techniques in optical communication have been applied to nonlinear impairment estimation, modulation format identification, and quality-of-transmission prediction [13]. These capabilities can improve assignment by allowing controllers to select channels that are more likely to meet performance targets. However, predictive models can be brittle when network conditions change or when training data no longer represent the operational environment. Intelligent assignment therefore requires uncertainty quantification, fallback rules, and human oversight. Governance mechanisms should define who is accountable when an automated assignment causes a service violation or disproportionately affects a customer group. The tussle among technical, commercial, and public interests in network design makes such accountability especially important [18].

Automation also raises questions of interoperability. If each vendor implements proprietary intelligence for spectrum assignment, multi-vendor networks may become fragmented. Standards and open interfaces are necessary to ensure that assignment decisions can be coordinated across equipment from different suppliers. Infrastructure studies emphasize that shared resources generate value through openness and interoperability, not only through proprietary optimization [19]. Intelligent assignment should therefore be designed with modularity, observability, and policy enforceability in mind. The goal is not to replace human decision-making entirely but to create a control environment in which automated decisions are transparent, reversible, and aligned with service and public objectives.

5. Robustness, Resilience, and Service Continuity

Optical networks are critical infrastructure, and their failure can disrupt finance, healthcare, government, education, and emergency response. Robustness in wavelength and frequency assignment means ensuring that service can continue or be restored quickly after fiber cuts, equipment failures, power outages, or correlated disasters. Resilience is broader than fault tolerance; it includes the capacity to absorb disturbances, adapt to changing conditions, and recover gracefully. Research on resilience and survivability in communication networks provides principles for designing systems that can withstand both expected and unexpected threats [16]. Network recovery practices for optical and multi-layer networks offer operational mechanisms such as protection switching, restoration, and pre-planned reconfiguration [17].

Intelligent assignment can improve resilience by reserving spectrum for protection paths, selecting routes with diverse physical risk, and adapting modulation formats to degraded links. However, resilience consumes resources. Dedicated protection reserves spectrum that may remain idle, while shared protection requires coordination and may fail under correlated failures. Restoration is more spectrum-efficient but slower and less predictable. The assignment problem must therefore balance capacity utilization against recovery time and service-level guarantees. In networks carrying deterministic or low-latency traffic, the tolerance for restoration delay may be very low, making protection more attractive despite its cost [12].

Resilience also has a spatial and social dimension. A fiber cut in a densely connected metropolitan area may affect many users but can often be rerouted through alternative paths. A cut in a rural or underserved region may sever the only available route, leaving communities isolated. Intelligent assignment should therefore consider geographic diversity and social criticality, not only network topology. Infrastructure planning that treats broadband as a shared resource can help align resilience investments with public interest [19]. Deployment histories in national broadband programs show that physical infrastructure decisions have long-lasting consequences for competition, inclusion, and economic development [20]. Wavelength and frequency assignment cannot solve these structural problems alone, but it can either mitigate or amplify them depending on how constraints and priorities are defined.

6. Sustainability and Energy Implications

The energy consumption of optical networks is a growing concern as traffic volumes increase and as networks expand into new regions. While optical transmission is generally more energy-efficient per bit than many alternative technologies, the absolute energy footprint of transponders, amplifiers, cooling systems, and data centers remains significant. Studies of green optical communications have examined energy limitations in transport and the potential

for efficiency improvements through technology and architecture [14]. Analyses of power trends in communication networks have shown that efficiency gains can be offset by demand growth, a dynamic sometimes described as rebound effects [15]. Intelligent wavelength and frequency assignment can contribute to sustainability by improving spectral efficiency, reducing regeneration, and minimizing unnecessary over-provisioning.

However, sustainability is not automatically achieved by maximizing spectral efficiency. Assignment decisions that increase transponder utilization may also increase cooling requirements or reduce equipment lifetime. Routing traffic through longer but less congested paths may save spectrum while consuming more amplifier power. Defragmentation can improve packing but may require temporary duplication of signals during reconfiguration. A life-cycle perspective is therefore necessary. Intelligent assignment should consider energy per delivered bit, embodied energy in equipment, and the operational energy associated with protection and restoration. It should also account for the interaction between optical and electronic layers, since grooming and regeneration can shift energy consumption between domains.

Sustainability also intersects with infrastructure policy. Public investments in broadband and optical backbone networks are often justified by economic development, education, and social inclusion. If those networks are not planned with energy constraints in mind, they may create long-term operational burdens for municipalities, cooperatives, and small operators. Intelligent assignment can support sustainability by making better use of existing spectrum and reducing the need for early equipment upgrades. Yet it cannot substitute for coherent planning of rights-of-way, conduit, power, and spectrum governance. Sustainable optical infrastructure requires coordination across engineering, policy, and finance.

7. Fairness, Governance, and Policy

Fairness in wavelength and frequency assignment is a relatively underdeveloped but increasingly important concern. Spectrum is a finite resource, and assignment decisions determine which services receive preferential treatment, which regions gain access to high-capacity connectivity, and which operators can compete effectively. In a commercial environment, fairness may be interpreted as non-discrimination among equivalent customers or as transparent rules for priority. In a public policy context, fairness may require attention to rural coverage, affordability, emergency services, and educational access. The design of assignment algorithms can embed these values through constraints, weights, or reserved capacity. Governance determines who sets those parameters and how disputes are resolved.

The governance of intelligent optical networks must address several tensions. First, automation may reduce human oversight, but accountability remains necessary. Second, centralized optimization may improve efficiency, but it may concentrate control in the hands of dominant operators or platform providers. Third, proprietary algorithms may create interoperability barriers, limiting competition and innovation. Fourth, data collection for machine learning raises privacy and security concerns, especially when telemetry reveals traffic patterns or customer behavior. Frameworks for understanding tussles in network design emphasize that technical architectures are shaped by competing interests and that governance must be built into the design process rather than added afterward [18]. Infrastructure theory further suggests that shared resources require institutional arrangements that balance private incentives with public value [19].

Policy implications extend to spectrum regulation, competition law, and universal service. If elastic optical networks become the primary transport infrastructure for digital services, then assignment policies may influence market entry and service quality. Regulators may need visibility into how spectrum is assigned in backbone and metropolitan networks, particularly where operators receive public rights-of-way or subsidies. At the same time, excessive regulation could slow innovation and investment. The challenge is to develop lightweight, transparent, and adaptable governance mechanisms that ensure fair access without micromanaging technical decisions. Broadband deployment experiences show that policy choices around infrastructure ownership, open access, and interconnection can shape outcomes for decades [20]. Intelligent wavelength and frequency assignment should therefore be understood as one component of a broader digital infrastructure governance regime.

8. Deployment Pathways and Institutional Considerations

Deployment of intelligent assignment capabilities will not occur uniformly. Incumbent operators with large legacy networks may adopt overlay solutions, gradually introducing flexible-grid equipment and software-defined control. New entrants may deploy greenfield elastic optical networks with greater architectural freedom. Research and education networks may serve as testbeds for advanced automation and deterministic services. Regional and rural networks may prioritize affordability, resilience, and simplicity over maximum spectral efficiency. These different pathways imply different requirements for intelligence. A nationwide backbone may benefit from centralized optimization and machine learning-based traffic forecasting, while a regional network may benefit from simple, rule-based assignment that local technicians can understand and maintain.

Institutional factors shape deployment as much as technical factors. Procurement practices determine whether open interfaces and multi-vendor interoperability are prioritized. Workforce training affects whether operators can manage automated reconfiguration and interpret telemetry. Vendor roadmaps influence the availability of bandwidth-variable transponders and flexible-grid reconfigurable optical add-drop multiplexers. Standards bodies and consortia shape the protocols that allow intelligence to function across domains. Financial models determine whether investments in defragmentation, monitoring, and spare capacity are justified. Intelligent assignment is therefore not a standalone innovation but a capability that depends on a wider ecosystem.

Cross-domain comparisons offer useful lessons. Cloud computing has demonstrated the value of software-defined resource orchestration, but it has also shown the risks of centralized control, opaque algorithms, and lock-in. Power grids have developed sophisticated approaches to real-time balancing, but they also illustrate the need for regulatory oversight and resilience against cascading failures. Transportation systems use dynamic pricing and routing, but they face equity concerns when optimization favors high-value users. Optical networks can learn from these domains by designing intelligence that is transparent, auditable, and aligned with public interest. The goal should be an assignment regime that is efficient enough to sustain growth, robust enough to protect critical services, and fair enough to maintain legitimacy.

9. Conclusion

Intelligent wavelength and frequency assignment in elastic optical networks is a system-level capability that connects physical-layer transmission, network control, operational practice, and public policy. The shift from fixed-grid to flexible-grid optical networking has created

new opportunities for spectral efficiency and service differentiation, but it has also introduced complex trade-offs involving fragmentation, resilience, energy consumption, fairness, and governance. Machine learning and automation can improve assignment decisions by predicting traffic, estimating quality of transmission, and managing failures, yet their benefits depend on data quality, interoperability, transparency, and accountability. Deterministic services and frequency-referenced multiplexing demonstrate that spectrum placement can support guaranteed bandwidth and low latency, but such guarantees must be balanced against broader network objectives.

The paper has argued that intelligent assignment should not be treated as a narrow optimization problem. It is a socio-technical practice embedded in infrastructures, institutions, and markets. Deployment pathways will vary across incumbents, new entrants, research networks, and regional operators, each with different constraints and priorities. Sustainability requires attention to energy per bit, equipment life cycles, and long-term operational burdens. Fairness requires explicit consideration of how spectrum access affects competition, inclusion, and service quality. Governance requires mechanisms for accountability, interoperability, and public oversight. Future research should therefore integrate algorithmic innovation with policy analysis, infrastructure studies, and participatory design. The ultimate measure of intelligence in optical spectrum assignment is not only how much capacity it extracts from fiber, but how well it supports resilient, sustainable, and equitable digital infrastructure.

References

1. Jinno, M., Takara, H., Kozicki, B., Tsukishima, Y., Sone, Y., & Matsuoka, S. (2009). Spectrum-efficient and scalable elastic optical path network: Architecture, benefits, and enabling technologies. *IEEE Communications Magazine*, 47(11), 66–73. <https://doi.org/10.1109/MCOM.2009.5307468>
2. Gerstel, O., Jinno, M., Lord, A., & Yoo, S. J. B. (2012). Elastic optical networking: A new dawn for the optical layer? *IEEE Communications Magazine*, 50(2), s12–s20. <https://doi.org/10.1109/MCOM.2012.6146481>
3. Klinkowski, M., & Walkowiak, K. (2011). Routing and spectrum assignment in spectrum sliced elastic optical path network. *IEEE Communications Letters*, 15(8), 884–886. <https://doi.org/10.1109/LCOMM.2011.060811.110281>
4. Christodoulopoulos, K., Tomkos, I., & Varvarigos, E. A. (2011). Elastic bandwidth allocation in flexible OFDM-based optical networks. *Journal of Lightwave Technology*, 29(9), 1354–1366. <https://doi.org/10.1109/JLT.2011.2125777>
5. Talebi, S., Alam, M. J., Katib, I., Khamis, M., Salama, R., & Rouskas, G. N. (2014). Spectrum management techniques for elastic optical networks: A survey. *Optical Switching and Networking*, 13, 34–48. <https://doi.org/10.1016/j.osn.2014.02.003>
6. Wang, Y., Cao, X., & Pan, Y. (2011). A study of the routing and spectrum allocation in spectrum-sliced elastic optical path networks. In *2011 IEEE International Conference on Communications (ICC)* (pp. 1–5). IEEE. <https://doi.org/10.1109/icc.2011.5963418>
7. Zhang, G., De Leenheer, M., Morea, A., & Mukherjee, B. (2013). A survey on OFDM-based elastic core optical networking. *IEEE Communications Surveys & Tutorials*, 15(1), 65–87. <https://doi.org/10.1109/SURV.2012.010912.00123>

8. Rouskas, G. N. (2016). Routing and spectrum allocation. In V. López & L. Velasco (Eds.), *Elastic optical networks: Architectures, technologies, and control* (pp. 77–103). Springer. https://doi.org/10.1007/978-3-319-30174-7_4
9. Chatterjee, B. C., Sarma, N., & Oki, E. (2018). Routing and spectrum allocation in elastic optical networks: A tutorial. *IEEE Communications Surveys & Tutorials*, 20(3), 2107–2125. <https://doi.org/10.1109/COMST.2018.2815397>
10. Musumeci, F., Rottondi, C., Nag, A., Macaluso, I., Zibar, D., Ruffini, M., & Tornatore, M. (2019). A tutorial on machine learning for failure management in optical networks. *Journal of Lightwave Technology*, 37(16), 4065–4075. <https://doi.org/10.1109/JLT.2019.2920329>
11. Mata, J., de Miguel, I., Durán, R. J., Merayo, N., Singh, S. K., Jukan, A., & Chamania, M. (2018). Artificial intelligence (AI) in optical networks: A review. *IEEE Communications Surveys & Tutorials*, 20(3), 2060–2089. <https://doi.org/10.1109/COMST.2018.2810658>
12. Zhou, Z., Wei, J., Luo, Y., Clark, K. A., Siljekens, E., Deakin, C., ... & Liu, Z. (2023). Communications with guaranteed bandwidth and low latency using frequency-referenced multiplexing. *Nature Electronics*, 6(9), 694–702.
13. Zibar, D., Piels, M., Jones, R., & Schäffer, C. G. (2016). Machine learning techniques in optical communication. *Journal of Lightwave Technology*, 34(6), 1442–1452. <https://doi.org/10.1109/JLT.2015.2504962>
14. Tucker, R. S. (2011). Green optical communications—Part I: Energy limitations in transport. *IEEE Journal of Selected Topics in Quantum Electronics*, 17(2), 245–260. <https://doi.org/10.1109/JSTQE.2010.2051216>
15. Kilper, D. C., Atkinson, G., Korotky, S. K., Goyal, S., Vetter, P., Suvakovic, D., & Blume, O. (2011). Power trends in communication networks. *IEEE Journal of Selected Topics in Quantum Electronics*, 17(2), 275–284. <https://doi.org/10.1109/JSTQE.2010.2050788>
16. Sterbenz, J. P. G., Hutchison, D., Çetinkaya, E. K., Jabbar, A., Rohrer, J. P., Schöller, M., & Smith, P. (2010). Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Computer Networks*, 54(8), 1245–1265. <https://doi.org/10.1016/j.comnet.2010.03.005>
17. Vasseur, J. P., & Pickavet, M. (2008). *Network recovery: Protection and restoration of optical, SONET-SDH, IP, and MPLS*. Morgan Kaufmann.
18. Clark, D. D., Wroclawski, J., Sollins, K. R., & Braden, R. (2005). Tussle in cyberspace: Defining tomorrow's Internet. *IEEE/ACM Transactions on Networking*, 13(3), 462–475. <https://doi.org/10.1109/TNET.2005.850223>
19. Frischmann, B. M. (2012). *Infrastructure: The social value of shared resources*. Oxford University Press.
20. Crawford, S. P. (2018). *Fiber: The coming tech revolution—and why America might miss it*. Yale University Press.