

Intelligent Modulation and Coding Adaptation for High-Speed Fiber-Optic Systems

Vilgot Holm

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

holm1977@ku.edu

Naomh Deady

Department of Computer Science, University of Alabama at Birmingham, Birmingham, AL, USA.

naomh451@uab.edu

Frado D. Bajwa

Department of Computer Science and Engineering, University of Nevada, Reno, Reno, NV, USA.

fradob@unr.edu

Rhiannon Colfer

School of Information Technology, University of Cincinnati, Cincinnati, OH, USA.

colfer1981@uc.edu

Abstract

High-speed fiber-optic systems are no longer best understood as static pipes optimized once at deployment. They are dynamic socio-technical infrastructures whose performance depends on the continuous alignment of modulation formats, coding rates, spectral provisioning, monitoring, and control. This paper examines intelligent modulation and coding adaptation as a system-level capability rather than a narrow physical-layer technique. It argues that adaptation must be designed across several layers: the optical transmission layer, the digital signal processing layer, the network control and management layer, and the governance layer that determines how capacity, latency, and resilience are allocated. The analysis considers trade-offs among spectral efficiency, reach, energy consumption, latency, hardware complexity, and operational robustness. It also explores how telemetry, machine learning, and elastic optical networking can support closed-loop adaptation while introducing new risks related to opacity, feedback instability, cybersecurity, and unequal service quality. Special attention is given to deployment constraints, infrastructure ownership, sustainability, fairness, and policy implications. The paper concludes that intelligent adaptation is most credible when it is treated as an accountable infrastructure practice: measurable, auditable, interoperable, and aligned with public-interest objectives such as equitable access, energy transparency, and resilient connectivity. The central challenge is not merely to select better modulation and coding parameters, but to build institutions and architectures capable of learning responsibly across heterogeneous high-speed optical environments.

Keywords

fiber-optic communications, modulation adaptation, coding adaptation, elastic optical networks, machine learning, infrastructure governance, sustainability, robustness, fairness.

1. Introduction

High-speed fiber-optic systems form the backbone of contemporary digital society, carrying data among data centers, metropolitan exchanges, submarine landing stations, research networks, and access aggregation points. Their engineering has traditionally emphasized capacity maximization under constraints of reach, bandwidth, and cost. Foundational analyses of fiber capacity established that optical transmission is bounded by complex interactions among noise, nonlinearity, spectral efficiency, and hardware limitations [1]. Subsequent work on advanced modulation formats showed that the choice of modulation is not merely a matter of transmitter design but a central determinant of network economics, reach, and upgradability [2]. Coherent detection and digital signal processing further transformed fiber systems by allowing phase, polarization, and amplitude to be exploited with unprecedented precision [3], [4]. In this setting, intelligent modulation and coding adaptation emerges as a response to a basic systems problem: the optimal transmission configuration is rarely stable across time, distance, traffic, fiber condition, or service obligation.

The need for adaptation intensifies as optical networks move from rigid wavelength provisioning toward more flexible, software-defined, and service-differentiated infrastructures. Power-efficient modulation formats, probabilistic shaping, and advanced forward error correction allow operators to trade off reach, rate, and margin in ways that were difficult to automate in earlier generations [5], [6]. At the same time, nonlinear propagation limits mean that simply increasing modulation order or launch power can degrade performance rather than improve it [7], [8]. Adaptation therefore becomes a problem of constrained optimization under uncertainty, where the system must decide not only what is technically possible but what is operationally prudent. The question is not whether modulation and coding can be adapted, but how adaptation should be governed when the underlying infrastructure is shared, long-lived, and economically consequential.

This paper develops a system-level account of intelligent modulation and coding adaptation in high-speed fiber-optic systems. It emphasizes architecture, trade-offs, deployment, governance, sustainability, robustness, fairness, and policy. Rather than isolating physical-layer algorithms, the discussion treats adaptation as an infrastructure capability that connects telemetry, analytics, control planes, service-level agreements, and public obligations. The argument proceeds through the layered architecture of adaptation, the trade-offs among modulation and coding choices, the role of sensing and closed-loop control, the promise and limits of artificial intelligence, the deployment and governance context, and the implications for robustness, fairness, sustainability, and policy. The conclusion synthesizes these themes and identifies conditions under which intelligent adaptation can be both effective and legitimate.

2. System Architecture and Adaptation Layers

Intelligent modulation and coding adaptation cannot be understood as a single control function. It operates across layered architectures that include the optical transmission layer, the digital signal processing layer, the network control and management layer, and the service orchestration layer. At the optical layer, adaptation may involve selecting modulation formats, coding rates, carrier frequencies, power levels, and spectral widths. At the digital layer, it may involve tuning equalization, forward error correction, constellation shaping, and impairment compensation. At the network layer, it may involve reconfiguring lightpaths, spectrum assignments, protection paths, and restoration policies. At the service layer, it may involve translating application requirements into optical performance targets. Each layer has its own

timescale, observability, and failure modes, which means that adaptation must be coordinated rather than centralized in a single controller.

Elastic optical networking provides an important architectural reference because it decouples service demand from fixed wavelength grids and allows spectrum to be allocated in variable-width slices [12], [14]. This flexibility makes adaptation more feasible, but it also increases the combinatorial complexity of resource assignment. A network that can change modulation and coding dynamically must also determine when changes are worthwhile, which flows should receive priority, and how guard bands, fragmentation, and restoration capacity should be managed. Frequency-referenced multiplexing illustrates a related direction in which guaranteed bandwidth and low latency are embedded more directly into the optical transmission architecture [13]. Such approaches suggest that adaptation is not only about maximizing throughput but also about making performance commitments stable enough for emerging applications.

The architectural challenge is compounded by the coexistence of legacy and new equipment. High-speed fiber systems are long-lived infrastructures with mixed generations of transponders, amplifiers, reconfigurable optical add-drop multiplexers, and control systems. An intelligent adaptation layer must therefore be backward compatible, vendor-neutral where possible, and capable of operating with incomplete telemetry. If adaptation depends on proprietary data models or closed interfaces, it risks creating lock-in and uneven capability across the network. Governance thus enters the architecture at an early stage: the decision to expose telemetry, standardize control interfaces, and document adaptation policies shapes who can participate in optimization and who bears the consequences of degraded performance.

3. Modulation, Coding, and Shaping Trade-Offs

The core trade-offs in modulation and coding adaptation are well known in principle but difficult to manage in practice. Higher-order modulation formats increase spectral efficiency by encoding more bits per symbol, but they also reduce tolerance to noise, nonlinearity, and residual impairments [2], [5]. Lower-order formats provide greater reach and robustness but consume more spectrum for a given data rate. Forward error correction can recover transmission errors, yet stronger coding typically adds overhead, latency, and computational burden [11]. Probabilistic constellation shaping offers finer granularity by adjusting the probability distribution of transmitted symbols, allowing the system to approach channel capacity more closely and to adapt rate and reach with greater precision [9], [10]. These techniques are not substitutes for one another; they form a design space in which adaptation must select a configuration that balances competing objectives.

The trade-offs are also economic and operational. A modulation format that maximizes capacity on a single link may increase energy consumption, cooling requirements, and hardware complexity across the network. A coding scheme that improves resilience may add latency that violates service guarantees for interactive or industrial applications. A shaping strategy that improves spectral efficiency may require more sophisticated digital signal processing and higher-speed electronics, shifting cost from fiber to transponders. In high-speed systems, the marginal gains from adaptation can be substantial, but they are often realized only when the entire system, including monitoring, margin allocation, and restoration, is designed to support them. Adaptation that ignores these couplings may optimize a local metric while degrading global service quality.

A further complication is that the optimal configuration is not stationary. Fiber aging, connector degradation, amplifier gain tilt, temperature variation, maintenance activity, and changing traffic patterns all alter the effective channel. Intelligent adaptation therefore requires continuous estimation of link conditions and service requirements. Yet excessive adaptation can itself be harmful: frequent modulation changes may cause transient errors, complicate fault diagnosis, and create instability in control loops. The design goal is not maximal adaptation but appropriate adaptation, where the rate and magnitude of changes are matched to the stability of the underlying infrastructure and the criticality of the services being carried.

4. Sensing, Telemetry, and Control Loops

Intelligent adaptation depends on observability. High-speed optical systems generate a wide range of telemetry, including optical power, amplifier gain, bit error rates, pre-forward-error-correction error rates, polarization states, spectral occupancy, latency, and transponder temperature. Digital coherent receivers provide rich signal information that can be used to infer channel impairments and distinguish among noise, nonlinearity, and filtering effects [4]. When this telemetry is aggregated across links and nodes, it becomes possible to construct a dynamic picture of network health and performance. However, telemetry is not self-interpreting. It must be calibrated, time-aligned, and contextualized with topology, service, and policy information. Without such contextualization, adaptation may respond to symptoms rather than causes.

Closed-loop control introduces both promise and risk. A control loop that detects degradation and adjusts modulation or coding can improve resilience and reduce manual intervention. But closed loops can also oscillate, overcorrect, or propagate local errors across the network. If multiple controllers act on overlapping resources without coordination, the result may be conflict and instability. The design of control loops must therefore include damping mechanisms, fallback modes, human oversight, and clear boundaries of authority. In safety-critical or high-value services, adaptation may need to be constrained by conservative guard bands and explicit approval workflows. The question is not whether automation is desirable, but how much autonomy should be delegated and under what accountability structures.

Telemetry also raises governance questions. Detailed performance data can reveal traffic patterns, customer behavior, and infrastructure vulnerabilities. If telemetry is shared across operators or with third-party analytics providers, privacy, commercial confidentiality, and cybersecurity must be addressed. Standardized, minimal, and purpose-limited data exchange may be necessary to enable multi-domain adaptation without creating excessive surveillance or dependency. In this sense, the sensing layer is not merely technical; it is a site where decisions about transparency, trust, and control are made.

5. Artificial Intelligence and Optimization

Artificial intelligence and machine learning have become prominent candidates for managing the complexity of optical adaptation. Surveys of AI methods in optical networks show applications ranging from fault detection and quality-of-transmission estimation to resource allocation and traffic prediction [15]. Machine learning can model nonlinear relationships that are difficult to capture with simple rules, and it can learn from historical telemetry to anticipate failures or optimize spectrum use [16]. In optical communication, learning techniques have been applied to impairment compensation, modulation classification, and receiver optimization [17]. These capabilities are attractive because high-speed systems

produce large data streams and operate under conditions that are too variable for static planning alone.

Yet the use of AI in infrastructure control must be approached with caution. Many machine learning models are opaque, data-dependent, and brittle outside their training distribution. An adaptive controller trained on one network topology may perform poorly on another with different amplifier spacing, fiber types, or traffic profiles. If the model is not interpretable, operators may struggle to diagnose failures or justify decisions to regulators and customers. AI systems can also amplify biases in data, privileging well-instrumented routes or high-value customers while neglecting older or rural infrastructure. The governance challenge is to ensure that AI-assisted adaptation improves overall network performance without creating hidden inequities or unaccountable failure modes.

A credible approach is to combine model-based reasoning with data-driven learning. Physics-informed models of fiber propagation, amplifier behavior, and noise accumulation can constrain the search space and provide explainable baselines [7]. Machine learning can then be used where it adds value, such as anomaly detection, forecasting, and fine-tuning of parameters. Hybrid architectures are more likely to be robust, auditable, and transferable across vendors and domains. They also support human oversight by exposing the assumptions and constraints under which adaptation occurs. The goal is not to replace optical engineering judgment but to augment it with faster, broader, and more systematic analysis.

6. Deployment, Infrastructure, and Governance

Deployment realities shape what intelligent adaptation can achieve. Fiber-optic systems are capital-intensive, geographically dispersed, and often governed by a mix of private operators, public agencies, consortia, and standards bodies. Upgrading modulation and coding capabilities may require new transponders, amplifiers, or monitoring equipment, and the benefits may accrue unevenly across the network. In competitive markets, operators may prioritize dense urban routes where returns are highest, leaving rural, remote, or disadvantaged areas with older and less adaptive infrastructure. Without deliberate policy intervention, intelligent adaptation could widen existing digital divides by making advanced performance available only where investment is commercially attractive.

Governance must therefore address the terms under which adaptation is deployed. Who decides when a link is reconfigured? How are service-level agreements updated? What recourse exists for customers whose performance degrades due to adaptation on shared infrastructure? These questions connect to broader debates about network neutrality and infrastructure as a shared social resource [21], [22]. If adaptation is used to prioritize certain traffic or customers, it must be transparent and consistent with regulatory obligations. If it is used to manage congestion, it should not become a pretext for discriminatory treatment. The technical flexibility of modulation and coding adaptation makes governance more, not less, important.

Infrastructure ownership and interconnection also matter. High-speed optical networks are rarely single-operator domains; they involve wholesale carriers, content providers, cloud platforms, submarine cable consortia, and research networks. Adaptation across domains requires trust, standardized interfaces, and clear liability rules. A failure caused by an adaptive change in one domain may affect many downstream services, raising questions about responsibility and compensation. Multi-domain adaptation is thus as much an institutional

problem as an engineering one. Success depends on agreements about data sharing, change notification, performance verification, and dispute resolution.

7. Robustness, Fairness, and Sustainability

Robust adaptation must be designed for failure, not just efficiency. High-speed fiber systems face fiber cuts, equipment failures, malicious attacks, and natural disasters. An adaptive system that continuously seeks higher spectral efficiency may reduce margins that are needed during emergencies. Conversely, a system that is too conservative may waste capacity and energy. The right balance depends on the criticality of the service, the availability of restoration paths, and the predictability of demand. Frequency-referenced multiplexing and related approaches that guarantee bandwidth and low latency point toward architectures in which determinism and adaptability are co-designed rather than treated as opposing goals [13]. Robustness also requires graceful degradation: when adaptation fails, the network should fall back to known-safe configurations rather than collapse.

Fairness is a less familiar but increasingly important dimension of optical adaptation. If machine learning or dynamic control allocates spectrum, coding, and restoration priority based on predicted revenue or customer value, it may systematically disadvantage public institutions, small enterprises, or marginalized communities. Fairness in infrastructure does not require identical service for all, but it does require that differential treatment be justified, transparent, and subject to review. Adaptation policies should include equity metrics alongside efficiency metrics, such as availability in underserved regions, latency for public services, and resilience for emergency communications. Without such metrics, optimization may be technically sophisticated but socially regressive.

Sustainability adds another constraint. The energy consumption of digital infrastructure is significant and growing, and the electricity intensity of data transmission depends on network design, equipment efficiency, and utilization [18], [19], [20]. Adaptive modulation and coding can reduce energy per bit by avoiding overprovisioning and by selecting configurations that match traffic demand. However, the additional processing required for adaptation, including sensing, machine learning, and high-speed digital signal processing, also consumes energy. The net effect is empirical and context-dependent. Sustainable adaptation requires lifecycle assessment, transparent reporting, and incentives for efficiency that do not simply shift environmental burdens from one part of the system to another.

8. Policy and Standards

Policy and standards are essential to making intelligent adaptation interoperable, accountable, and fair. Standards bodies can define common telemetry models, control interfaces, and performance metrics that allow multi-vendor and multi-domain coordination. They can also specify minimum safety and fallback requirements so that adaptive systems do not create opaque or unpredictable behavior. Regulators, meanwhile, can require transparency about how adaptation affects service quality, especially where networks receive public subsidies or operate as essential infrastructure. Policy should encourage innovation while preventing adaptation from becoming a mechanism for hidden discrimination or unaccountable degradation.

A key policy tension is between flexibility and predictability. Operators want the freedom to optimize spectrum, modulation, and coding in response to changing conditions. Customers, application providers, and regulators want stable performance and clear commitments. This tension can be managed through tiered service agreements, published adaptation policies, and

independent auditing. For example, an operator might disclose that certain premium services receive priority restoration or wider margins, while basic services receive best-effort adaptation. Such disclosure does not eliminate trade-offs, but it makes them visible and contestable. In democratic infrastructure governance, visibility is a precondition for legitimacy.

International coordination is also necessary because high-speed fiber networks cross borders through submarine cables and terrestrial backbones. Divergent national rules on data sharing, cybersecurity, and service priority can impede adaptation or fragment the global internet. Common principles, such as proportionality, non-discrimination, security by design, and environmental transparency, can provide a baseline while allowing local implementation. The goal is not a single global adaptation regime, but a layered governance ecosystem in which technical standards, market rules, and public-interest obligations reinforce one another.

9. Conclusion

Intelligent modulation and coding adaptation is a defining capability for high-speed fiber-optic systems, but it is not merely a matter of selecting better transmission parameters. It is an architectural, operational, and governance challenge that spans physical-layer design, digital signal processing, network control, service orchestration, and public policy. The most promising systems will combine elastic optical networking, rich telemetry, hybrid model-based and data-driven analytics, and careful control-loop design [12], [14], [15], [17]. They will also recognize that adaptation creates winners and losers, consumes energy, and introduces new forms of complexity and dependency.

The central argument of this paper is that intelligent adaptation should be treated as an accountable infrastructure practice. Effectiveness requires accurate sensing, robust optimization, and graceful fallback. Legitimacy requires transparency, fairness, interoperability, and policy oversight. Sustainability requires lifecycle thinking and energy transparency, not simply local efficiency gains [18], [19], [20]. Robustness requires margins, diversity, and human oversight, especially when adaptive systems control critical capacity and latency commitments [13]. Fairness requires that optimization metrics include underserved regions, public services, and emergency communications, not only commercial value.

Future research should examine how adaptation policies perform across heterogeneous topologies, how multi-domain governance can be implemented without excessive centralization, and how energy and equity impacts can be measured reliably. It should also investigate how operators, regulators, and communities can co-design adaptation regimes that are both innovative and trustworthy. High-speed fiber-optic systems will remain foundational to digital economies and societies. The question is whether their intelligence will be deployed as a narrow efficiency tool or as a contribution to resilient, sustainable, and equitable communication infrastructure. The latter requires not only better algorithms but also better institutions.

References

1. Essiambre, R. J., Kramer, G., Winzer, P. J., Foschini, G. J., & Goebel, B. (2010). Capacity limits of optical fiber networks. *Journal of Lightwave Technology*, 28(4), 662-701. <https://doi.org/10.1109/JLT.2009.2039464>
2. Winzer, P. J., & Essiambre, R. J. (2006). Advanced optical modulation formats. *Proceedings of the IEEE*, 94(5), 952-985. <https://doi.org/10.1109/JPROC.2006.873438>

3. Kikuchi, K. (2016). Fundamentals of coherent optical fiber communications. *Journal of Lightwave Technology*, 34(1), 157-179. <https://doi.org/10.1109/JLT.2015.2463719>
4. Savory, S. J. (2010). Digital coherent optical receivers: Algorithms and subsystems. *IEEE Journal of Selected Topics in Quantum Electronics*, 16(5), 1164-1179. <https://doi.org/10.1109/JSTQE.2010.2044751>
5. Agrell, E., & Karlsson, M. (2009). Power-efficient modulation formats in coherent transmission systems. *Journal of Lightwave Technology*, 27(22), 5115-5126. <https://doi.org/10.1109/JLT.2009.2029064>
6. Kahn, J. M., & Ho, K. P. (2004). Spectral efficiency limits and modulation/detection techniques for DWDM systems. *IEEE Journal of Selected Topics in Quantum Electronics*, 10(2), 259-272. <https://doi.org/10.1109/JSTQE.2004.826575>
7. Poggiolini, P. (2012). The GN model of non-linear propagation in uncompensated coherent optical systems. *Journal of Lightwave Technology*, 30(24), 3857-3879. <https://doi.org/10.1109/JLT.2012.2217729>
8. Ellis, A. D., Zhao, J., & Cotter, D. (2010). Approaching the non-linear Shannon limit. *Journal of Lightwave Technology*, 28(4), 423-433. <https://doi.org/10.1109/JLT.2009.2030693>
9. Buchali, F., Steiner, F., Böcherer, G., Schmalen, L., Schulte, P., & Idler, W. (2016). Rate adaptation and reach increase by probabilistically shaped 64-QAM: An experimental demonstration. *Journal of Lightwave Technology*, 34(7), 1599-1609. <https://doi.org/10.1109/JLT.2016.2535191>
10. Fehenberger, T., Alvarado, A., Böcherer, G., & Hanik, N. (2016). On probabilistic shaping of quadrature amplitude modulation for the nonlinear fiber channel. *Journal of Lightwave Technology*, 34(21), 5063-5073. <https://doi.org/10.1109/JLT.2016.2594271>
11. Kschischang, F. R., & Steiner, F. (2019). Probabilistic shaping and forward error correction. *Journal of Lightwave Technology*, 37(1), 52-60. <https://doi.org/10.1109/JLT.2018.2878325>
12. 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>
13. 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.
14. 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>
15. Mata, J., de Miguel, I., Durán, R. J., Merayo, N., Singh, S. K., Jukan, A., & Chamania, M. (2018). Artificial intelligence (AI) methods in optical networks: A comprehensive survey. *Optical Switching and Networking*, 28, 43-57. <https://doi.org/10.1016/j.osn.2017.12.006>

16. Khan, F. N., Fan, Q., Lu, C., & Lau, A. P. T. (2019). An optical communication's perspective on machine learning and its applications. *Journal of Lightwave Technology*, 37(2), 493-516. <https://doi.org/10.1109/JLT.2018.2872562>
17. 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.2015.2508503>
18. Lange, S., Pohl, J., & Santarius, T. (2020). Digitalization and energy consumption. Does ICT reduce energy demand? *Ecological Economics*, 176, 106760. <https://doi.org/10.1016/j.ecolecon.2020.106760>
19. Aslan, J., Mayers, K., Koomey, J. G., & France, C. (2018). Electricity intensity of Internet data transmission: Untangling the estimates. *Journal of Industrial Ecology*, 22(4), 785-798. <https://doi.org/10.1111/jiec.12630>
20. Coroamă, V. C., & Hilty, L. M. (2014). Energy consumption of information and communication technology (ICT) in Germany up to 2020. In L. M. Hilty & B. Aebischer (Eds.), *ICT innovations for sustainability* (pp. 131-146). Springer. https://doi.org/10.1007/978-3-319-09228-7_8
21. Wu, T. (2003). Network neutrality, broadband discrimination. *Journal of Telecommunications and High Technology Law*, 2, 141-179.
22. Frischmann, B. M. (2012). *Infrastructure: The social value of shared resources*. Oxford University Press.