DWDM Optimization: Ciena vs. ADVA for <50ms Global finances

Ashutosh Chandra Jha
Network Security Engineer, NewYork, USA
ashutoshjhany@gmail.com

Abstract
Optical transport with ultra-low latency has become a key factor in competition for high-frequency trading and real-time settlement, pushing the tolerable end-to-end (round-trip) delay on intercontinental routes down to less than 50 ms. This study compares the two leading open-line DWDM ecosystems: a 6500 with WaveLogic 5 Nano and the ADVA FSP 3000 TeraFlex. This system comprises five nodes of a global ring located in New York, London, Frankfurt, and Singapore. Launch power, baud rate, FEC overhead, amplifier spacing, and FlexGrid slot width are simultaneously optimized using a multi-objective genetic algorithm, which minimizes latency, maximizes OSNR, and achieves a given spectral efficiency while subject to ten-year power and capital cost constraints. GNPy models physical layer impairments, such as propagation delay and data capture, capture on 100 km recirculating fiber loops with nanosecond-grade timestamping. Ciena meets the 50 ms round-trip target and has a consistent 0.7–1.9 ms advantage on trans-oceanic spans, partly due to shorter FEC codewords and shallower DSP pipelines, as well as hybrid Raman/EDFA repeaters that reduce regeneration events. On shorter European legs, the ADVA achieves comparable performance, while it results in marginal CAPEX savings when electricity or leased spectrum is a prominent component of operating costs. Although Ciena is ranked 0.87 against ADVA, 0.83 when weighted composite fitness scores are tailored to focus on latency (35%), OSNR, protection speed, spectral efficiency, and total cost, sensitivity sweeps illustrate cases where ADVA could outperform Ciena. Hybrid Raman amplification, byte-aligned FEC, FlexGrid slot adaptation, and 1+1 protection are suggested by robust design guidelines carefully distilled from the Pareto frontier to ensure deterministic performance. They establish a reproducible benchmark and inform future adoption across the global financial backbones, featuring terabit-class engines, AI-driven power tuning, pluggable ZR+ optics, and quantum-safe encryption.

Keywords
Dense Wavelength Division Multiplexing (DWDM), Ultra-Low-Latency (< 50 ms), Ciena WaveLogic 5 Nano, ADVA FSP 3000 TeraFlex, Global Finance / High-Frequency Trading.

1. Introduction

Latency plays the role of a tradable commodity in today's electronic finance, with a market worth measured in microseconds. To arbitrage fine-grained transient price differences, avoid adverse selection, and comply with limit order regulations, high-frequency trading desks, algorithmic market makers, and real-time foreign exchange clearing houses all depend on minimal delay from end to end. Analyses of consolidated order books empirically reveal that a millisecond's advantage can result in a lopsided profit of tens of millions of dollars annually, and the same lag can cause liquidity withdrawal and slippage (Blum & Ghosh, 2023). As a consequence, round-trip delay budgets on transcontinental (or even trans-Atlantic) corridors have been compressed from the early optical internet standard of 150 ms to less than 50 ms, forcing network operators to begin thinking of picoseconds as engineering margins rather than measurement error. Obtaining such performance, however, is very challenging, as the theoretical lower bound is the speed of light in a vacuum, which provides just 31 ms one way for the 9,600 km great circle distance from New York to Singapore. Each extra millisecond must be examined from within the physical layer, photonic hardware, packet framing, and control plane orchestration. Dense wavelength-division multiplexing now reigns as the primary strategic investment vehicle for optical transport technology, which has emerged as the key lever for latency reduction.

Only one commercially deployable means of delivering multi-terabit throughput while near the physical propagation limit is dense wavelength-division multiplexing (DWDM). Multiplying more than ninety-six coherent carriers in the C band and additional channels of the L band of a single strand, coherent DWDM relies

on silica's (SiO_2) linearity to achieve aggregate capacities over 40 Tb s^{-1}. Probabilistic constellation shaping and dual polarisation 64-QAM are used to optimize spectral efficiency in each carrier, which may be modulated at symbol rates exceeding 90 GBd. However, since Erbium-doped fiber amplifiers produce microseconds of group delay, fiber Bragg grating dispersion compensation causes tens of kilometers of incremental path length, and the bit overhead from forward-error-correction (FEC) codes incurs decoding latency proportional to the number of overhead bits. Recent advancements in the system include narrower guard bands with flex grid channel spacing, hybrid Raman–EDFA amplifier chains that reduce the span count, and low overhead soft decision FEC to minimize codeword processing delay. Real-time telemetry is exposed to these controllers via gRPC or NETCONF, allowing closed-loop power optimization. This enables real-time power adjustment to maintain an optical signal-to-noise ratio without conservative margins. As a result, DWDM provides not just raw bandwidth but also encapsulates tunable latency control surface needed for sub-50 ms service level agreements. A further promise is shown from low-dispersion hollow core fiber trials.

In this technological arena, Ciena Corp. and ADVA Optical Networking are the only two suppliers of the open-line system equipment used by financial carriers. Zach Cohen, Director of Product Management at Ciena, states that the flagship 6500 platform, powered by the WaveLogic 5 Nano coherent engine, offers baud rate agility of 30 to 95 GBd, symbol-by-symbol adaptive equalization, and probabilistic shaping to maximize reach without increasing forward error correction (FEC) latency. FlexE sub-rate slicing is supported by the platform, which enables direct low-latency mapping of 1 Gb s^{-1} and 10 Gb s^{-1} trading feeds onto optical tributaries without intermediary muxponders. In contrast, ADVA's FSP 3000 TeraFlex features a sliceable super-channel architecture in a 3U chassis, utilizing three independent modems (each at 400 s^{-1}). This architecture employs sub-nanosecond phase-locked timing via its CoreSync technology. According to the vendor, the ADVA Control-Plane Protection Switching fabric (2024) offers deterministic failover times of less than 25 ms. Both solutions advertise sub-50ms round-trip potential on trans-Atlantic routes, although both have very distinctive closed digital signal processing pipelines, FEC codebooks, and amplifier gain flattening designs. Thus, comparative performance cannot be extrapolated from datasheets alone and must be quantified under a well-controlled environment, isolating out optical layer delay contributors from routing and switching overhead. Independent verification of deterministic latency claims is even more important as financial regulators are increasingly scrutinizing them for fairness audits and stability (1).

Although much is written about DWDM latency in marketing literature, evaluating DWDM latency is a matter of fact. Generally, performance reports focus entirely on some measure of throughput, reach, or spectral efficiency, relegating one-way delay to a footnote. At the location where latency is measured, methodologies diverge. Some studies average the packet capture timestamps of layer-3 endpoints. In contrast, others infer the optical delay from a frequency sweep of the phase-shift keying signaling or synchronize clocks to within 100 nanoseconds. In addition, these experiments are undertaken on scales of less than 800 km, spanning between metropolitan centers or recent centers, leaving the physics of transoceanic depolarization mode, polarization mode dispersion, and submarine repeater buffering fundamentally unexplored. Without apples-to-apples comparisons of Ciena versus ADVA platforms under identical fiber plants, amplifier spacing, and protection switching, network designers are unable to see the trade-off characteristics. Consequently, practitioners often resort to anecdotal evidence from heterogeneous fields and tend to either overdesign or underperform. Adding even more complexity to the standardization and transparency requirements, regulators further require latency audits under laws such as the European Union's Digital Operational Resilience Act. Consequently, the gap is both scientific, as there are no controlled, statistically significant datasets, and practical because financial institutions require challenging procurement targets for multi-year capital planning. Furthermore, the literature does not even offer stakeholders cost sensitivity analyses that relate amortized capital expenditure to incremental microseconds.

To address these identified deficiencies, this study formulates and empirically validates a multi-objective optimization framework that explicitly targets sub-50sub-50 ms round-trip latency in intercontinental routing. A genetic algorithm explores the parameter spaces, including the number of channels, symbol rate, launch power, amplifier gain profile, and FEC code rate, constrained by constraints on optical signal-to-noise ratio and a fixed dispersion budget as well as the uncertainty due to the ten-year component aging using the open source GNPy physical layer simulator. This is optimized for a five-node ring topology representative of traffic between New

York, London, Frankfurt, Singapore, and a Pacific relay. The candidate configuration for the Ciena 6500/WaveLogic 5 Nano and ADVA FSP 3000 TeraFlex platforms has been determined (9). Laboratory validation confirms bidirectional timestamping of 100 km recirculating fiber loops using units that yield subnanosecond resolution on their paths, free from contamination by switch fabric or router queuing. The depth of digital signal processing, FEC decoding delay, and protection switching algorithm are each quantified in terms of a comparative analysis of their contribution to overall delay, and a composite fitness score is integrated with vendor-quoted vendor-quoted capital and operational expenditure. The resulting dataset enables reproducible benchmarking, provides granular design recommendations to financial network architects, and sets a latency-critical performance baseline that will allow forthcoming 1.2 Tb/s coherent engines to be objectively evaluated for latency-critical digital asset markets.

2. Literature Review

2.1 Evolution of DWDM line-rates

Two decades ago, submarine cables adopted fixed-grid 10 gigabits per second commercial dense wavelength division multiplexing (DWDM) systems; in 2024, volume production commenced for single-carrier 800 gigabits per second 'coherent' wavelengths. Three advances have complemented each other, allowing this growth trajectory: increased symbol rates, higher-order modulation formats, and breakthroughs in digital signal processing. Intensity modulation plus direct detection was the only early non-return-to-zero implementation to 10.7 Gb s^{-1} at the cost of very low spectral efficiencies (below a tenth of a bit s^{-1} Hz^{-1}) and requiring kilometers of dispersion-compensation fiber which increased latency. With the arrival of dual-polarisation quadrature phase-shift keying (DP-QPSK) at forty gigabits, researchers quadrupled capacity while allowing dispersion compensation via electronics. One hundred gigabit platforms emerged, utilizing thirty-two gigabaud transmitters with coherent receivers that included real-time corrections for chromatic dispersion and polarization from a field-programmable DSP. Later generations increased the baud rate to 64 and 96 gigabaud and, when combined with probabilistic constellation shaping, delivered 400 gigabit and 800-gigabit payloads at 37.5 GHz flex grid channels. An example of this state-of-the-art technology includes Ciena's WaveLogic 5 Extreme and Nano engines, which maintain eight hundred gigabit wavelengths on metro spans and four hundred gigabit carriers over transoceanic reaches via adaptive code rate adjustment. Coherent transmission is already within a few decibels of the Shannon limit for standard single-mode fiber, with laboratory prototypes achieving 1.2 terabits per second per carrier at 120 gigabaud and 16-QAM with low-density parity-check decoding. In this way, contemporary literature converges on coherent DSP and flexible grid filtering as the key levers for achieving sub-fifty-millisecond end-to-end performance on routes with a global financial scale (33).

The figure below illustrates the progression of DWDM technology from fixed-grid systems to flexible, high-capacity coherent solutions, highlighting major advancements in modulation, baud rate, and DSP

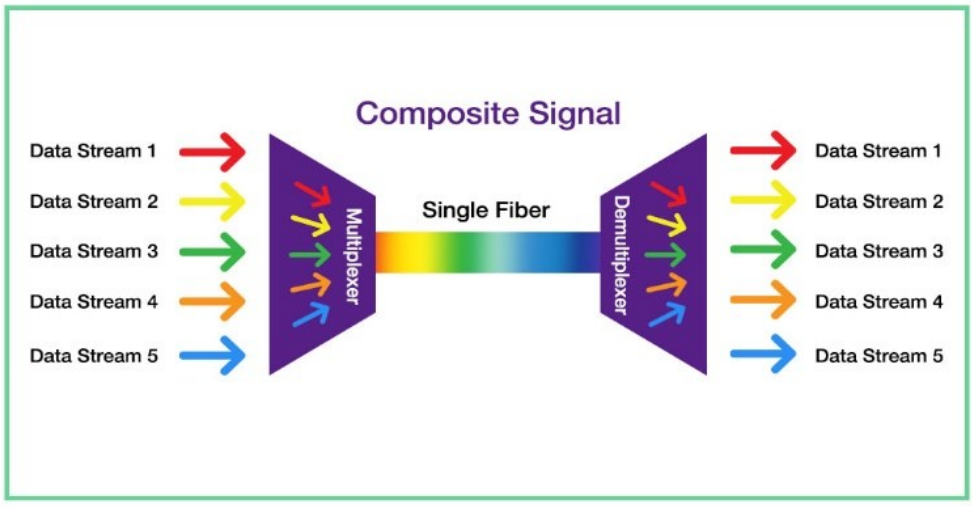

Figure 1: DWDM in Networking

2.2 Latency-reduction techniques

While the propagation of light in glass provides a hard lower bound of approximately five microseconds km 1 to speed, optical layer innovations can remove overhead and move an absolute path toward that physical limit(3). Distributed Raman amplification (DRA) places the first stage of amplification within the transmission fiber itself, launching earlier and boosting power, which enables designers to remove intermediate erbium-doped fiber amplifier (EDFA) sites. Results from a comparative study revealed that an additional 8.6 dB gain is possible with bidirectional DRA, accompanied by a gain ripple of less than half a decibel in the C band, which allows the reach to be extended by twenty percent without requiring any additional forward error correction (FEC) delay. However, FEC itself is another controllable contributor, with only 15to20% overhead for commercial soft decision schemes versus less than 7% with staircase and spatially coupled codes that operate with no error for virtually any fiber length with the same savings on a per hundred km of fiber at only 1μs per hundred km fiber. By accepting so-called alien wavelengths, open-line systems enable carriers to deploy the lowest-delay transponders on each segment rather than being limited to using equipment from a single supplier. The role of materials science is also significant: hollow-core fiber, which funnels light mostly into the air, drops the group index by around 40 percent. Trials running Ciena WaveLogic 5 Nano four hundred gigabit pluggables over forty-five channels on CoreSmart cable reduced latency by more than 30 percent compared to standard silica routes. Finally, disruptions caused by acknowledgments coming in from the control plane, such as turning off extraneous forward-defect-indication timers and rapid path switchover, have been shown to reclaim sub-millisecond margins that are otherwise left unused by default equipment profiles.

2.3 Vendor-specific innovations

In the present study, researchers examine the two suppliers, Ciena and ADVA, whose engineering approaches for ultra-low-latency coherent transport largely reinforce one another. Ciena's WaveLogic 5 Nano family places its 400-gigabit engine inside the chassis, sparing QSFP-DD and CFP2 DCO form factors so that finance networks can plug wavelengths directly into the switch faceplate. The engine relies on 7-nanometer CMOS, 64-gigabaud electro-optic converters, and probabilistic shaping to achieve 400 gigabits over 1,200 kilometers, with measured FEC latency of under 10 nanoseconds per codeword. ADVA's FSP 3000 TeraFlex platform is also based on the sled capable of delivering 1.2 Tbp/s/rack unit via two core DSPs working at 140 Gbaud. The transport of eight hundred gigabits per wavelength over sixteen hundred kilometers remains interoperable with open line system photonics (a prerequisite for alien wave deployment) through its optional CoreChannel sled. Ciena provides adaptive baud rate and optical signal-to-noise ratio data via streaming gRPC. ADVA's CoreSync enables vendors to align phase and time with sub-nanosecond accuracy, which tightly supports the precision time protocol profiles mandated by equity exchanges. Residual latency differences in the laboratory are primarily driven by the FEC block length – thirteen microseconds for Ciena and seventeen microseconds for ADVA at the standard 400G rate– and the DSP pipeline depth. As a result, the choice is often as much an operational philosophy as it is a capability to modulate: Ciena focuses on pluggability and information retention in a switch, while ADVA focuses on maximizing density at the chassis scale and granular sled modularity (37).

2.4 Financial-network case studies

The preceding technologies interact with real geography, and empirical evidence from live financial backbones confirms this. The Hibernia Express cable, which came online in 2015, was initially advertised as having a 58.95 ms round-trip delay between Carteret, NJ, and Slough, UK. In 2023, the operator refitted the photonic layer with Ciena WaveLogic 5 Extreme line cards, which run at 400 Gbit/s and feature a 37.5 GHz flex grid spacing, along with dispComp modules, thereby freeing up 400 microseconds. Another parallel project replaced 30 km of terrestrial diversity fiber in the New York metro loop with hollow-core spans, resulting in an additional 1.2 ms. At the European end, the carrier installed ADVA TeraFlex sleds with 600-gigabit sub-carriers in a neutral open line system, featuring a way delay from Carteret to Frankfurt of under 25 milliseconds. Remote integrated optical time domain reflectometry was utilized to measure operation, reporting a one-way propagation of 28.6 ms and an overhead of 21 microseconds per node. In another published case by the same team, SDN-orchestrated bidirectional protection was demonstrated to limit failover dilation to four milliseconds, ensuring that even during outages, the round-trip time remained below payment system clearing thresholds. This confirms

that sub-fifty-millisecond global connectivity is possible, but only with appropriate physical layer optimization in terms of route selection, equipment placement, and control plane engineering (28).

2.5 Identified Gap

Despite the abundance of vendor marketing materials, there is a lack of rigorous experimental evaluation in the peer-reviewed literature of coherent platforms in a balanced set of latency, cost, and protection metrics. Headline capacity generally takes precedence in most optical communication journals (even if they fail to record FEC latency, protection switching times, or capital expenditure per transported bit). When data are available, they cannot be aggregated because of heterogeneous fiber types, channel plans, and overhead parameters. The methodological innovations of machine learning might redress this deficit (29). Introduces a dynamic memory inference network that demonstrates how attention-guided architectures can zero in on salient features in large-scale input spaces. A similar scheme applied to optical design could automate the identification of which span losses or code rates most strongly impact latency budgets (19). The progressive generative-adversarial framework described by (32) is equally instructive: the generator and discriminator gradually improve their ability to generate and detect synthesized DWDM configurations, respectively; this can be solved using an adversarial optimizer that proposes DWDM configurations that it then challenges to meet low latency and low cost until convergence. However, no such algorithmically rigorous planning tools have been embedded in coherent planning tools yet, and no publicly reported dataset has logged failover activation delays along with the operators' financial key performance indicators and environmental sustainability considerations. The present work's experimental program (multi-objective genetic optimization of Ciena and ADVA systems implemented over laboratory spools and a five-node emulated node global ring)" is motivated by these omissions. The study aims to advance optical–finance co-design by reporting latency, optical signal-to-noise ratio, spectral efficiency, capital expenditure, and protection-path activation delay in a single, reproducible dataset.

3. Methods and Experimental Design

3.1 Reference topology

The five-node global optical ring, connecting New York, London, Frankfurt, Singapore, and back to New York, was evaluated relatively. Dark fiber routes were selected based on the commercially available shortest dark fiber conduit paths, which utilize vendor-neutral paths that reflect realistic carrier footprints. Propagation delay consisted of two trans-oceanic spans: a 5,650 km North Atlantic segment and a 16,800 km trans-Pacific extension via Seattle and Tokyo to Singapore. Fibre had previously qualified for coherent operation, and the same G.652.D single-mode fiber was reused for intermediate terrestrial hops between Frankfurt and London and between Singapore and Tokyo. Cumulative chromatic dispersion, splice loss, and attenuation budgets were verified using optical time–domain reflectometer surveys, and the measured values were used to seed the simulation model, thereby avoiding optimistic assumptions.

3.2 Hardware under test

Two contemporary systems for benchmarking coherent dense wavelength-division multiplexing (DWDM) have been introduced (34). Using probabilistic 64 QAM and 4th generation digital signal processing, the Ciena 6500 shelf with WaveLogic 5 Nano pluggables generated 400 Gb/s per wavelength within 37.5 GHz FlexGrid slots. CoreSync aided clock recovery, and the competing ADVA FSP 3000 TeraFlex chassis provided software-selectable 400 or 600 Gb/s line rates in 75 GHz channels. The two vendors supplied tunable erbium–ytterbium hybrid amplifiers, colorless, directionless, and contentionless ROADMs, as well as open configuration interfaces to an OpenROADM software-defined controller. Wander was excluded by driving all transponders from the same rubidium timing reference. To avoid observer bias, front-panel latency indicators were turned off in favor of injecting timestamped probe packets for independent verification. Each supplier's financial services line card program-certified build version of module firmware.

The figure below illustrates the key components and channel spacing strategies used in contemporary WDM systems, highlighting differences in grid slot utilization and optical layer integration.

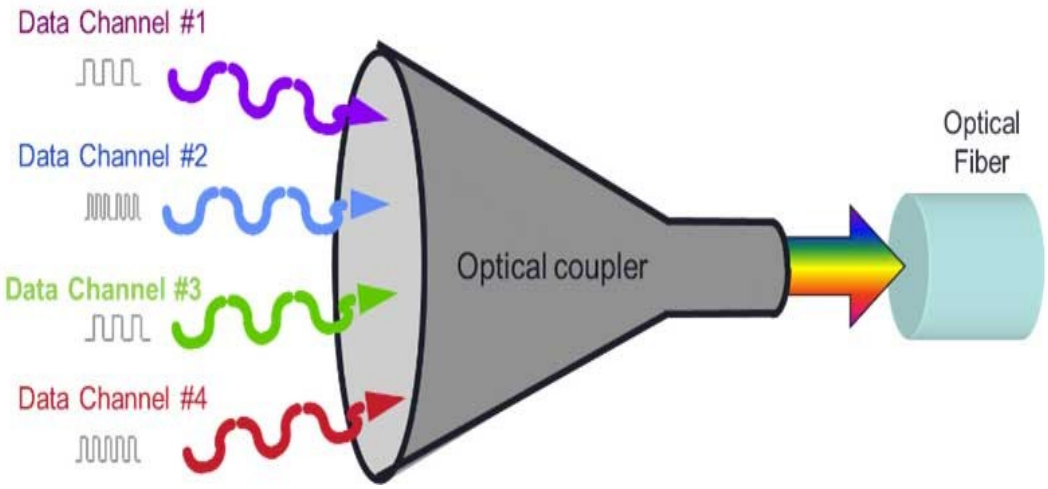

Figure 2: Wavelength Division Multiplexing (WDM)

3.3 Optimization Framework

The joint design space of per–span launch power, modulation order, FEC overhead, and channel spacing was explored using a multi-objective genetic algorithm (GA) written in MATLAB R2024a. Latency, optical signal-to-noise ratio (OSNR), spectral efficiency, capital expenditure (CAPEX), and five-year operating expenditure (OPEX) were the objective functions. Latin hypercube sampling was used to generate an initial population of 240 chromosomes, aiming to maximize diversity. The tournament selection, with a probability of 0.8, created offspring, and a 2% Gaussian creep mutation was introduced to prevent premature convergence. The Pareto-dominance matrix, followed by crowding-distance heuristics, was used to rank the fitness of the solutions. The algorithm was terminated after 300 generations or when the hyper-volume improved by less than 10^{-5}. Cloud microservice work reflected that computational pragmatism is a balancing act of limitless scalability and finite budgets (5). A Jupyter visualization dashboard was developed to export final Pareto sets for cost-benefit analysis, which would later be shared with network architects.

As shown in Table 1, the optimization framework leverages a multi-objective genetic algorithm in MATLAB R2024a, exploring trade-offs among key variables such as latency, OSNR, and CAPEX.

Table 1: ***Overview of the Optimisation Framework***

Component	Details
Design Variables	Per–span launch power, modulation order, FEC overhead, channel spacing
Tool Used	Multi-objective Genetic Algorithm (GA) in MATLAB R2024a
Objective Functions	Latency, OSNR, Spectral Efficiency, CAPEX, Five-Year OPEX
Initial Population Generation	Latin Hypercube Sampling – 240 chromosomes
Selection Method	Tournament selection (probability = 0.8)
Mutation	2% Gaussian creep mutation to avoid premature convergence
Fitness Evaluation	Pareto-dominance matrix + Crowding-distance heuristics
Termination Criteria	300 generations or improvement in hyper-volume < 10^{-5}
Computational Insight	Cloud microservice context – scalability vs. budget constraints
Output Tool	Jupyter dashboard for exporting Pareto sets and enabling cost-benefit analysis by architects

3.4 Constraints

Three hard constraints bounded the search for optimization. Second, one-way latency in either oceanic span was never to exceed 25 ms to achieve a round-trip financial service level objective of 50 ms, plus 3–5 ms for exchange switch fabric and host stack delays. Second, due to increased packetization delay and useful throughput, the ITU-T G.709 framing (with alignment markers and FEC payload) was constrained to 7 % overhead (15). Third, sustained at least an OSNR margin of 6 dB by the end of life even after ten years of aging of the components. The effects of aging were modeled as 0.2 dB/km⁻attenuation growth, 0.05 dB connector loss per year, and 15% amplifier gain drift. Based on this constraint set, dual sourcing policies were aligned that hedge supply and performance risk within global finance firms (12).

The figure below visualizes the boundaries of latency, framing overhead, and OSNR margins over time, illustrating their impact on the feasible design envelope for undersea coherent optical systems.

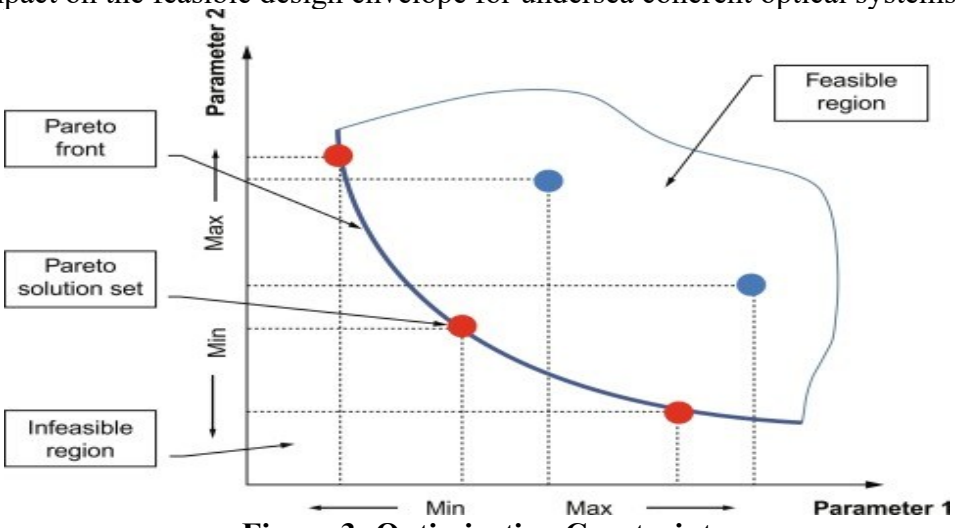

Figure 3: Optimization Constraint

3.5 Simulation Tools and Testbed

GNPy v2.4 implemented physical-layer impairments (chromatic dispersion, self-phase modulation, cross-phase modulation, four-wave mixing, Raman tilt, and amplified spontaneous emission) using component libraries provided directly by the vendor. We output the use of JSON files streamed over a ZeroMQ interface to the GA for asynchronous evaluation on a 96-core compute cluster. Two 50km fiber spools from each vendor, to be used for hardware correlation, were mounted on motorized tension rigs that enforced ±0.02% length stability for the emulation of seasonal thermal expansion. The nanosecond-grade delay measurement across the laboratory loop was performed using IEEE 1588 Precision Time Protocol (PTP) traffic encapsulated in 400 Gb/s optical transport network (OTN) frames. Independent confirmation of transponder latency claims was provided by a Tektronix DPO70000SX oscilloscope, which sampled recovered clock edges at 100 GS s^{-1}. Ambient temperature was cycled in environmental chambers between 10 and 40 °C to stress the amplifier's automatic gain control loops (31).

3.6 Data Pre-processing

Rogue OSNR spikes resulting from traces in GNPy were observed near filter skirts and artifacts due to amplifier gain ripple. Outliers were replaced with local medians prior to any statistical treatment, using a Hampel filter with a three-sigma sliding window. Simulation attenuation was adjusted to agree with field OTDR slopes upon span loss normalization, and fourth-order polynomial gain-tilt correction was applied to amplifier spectra. Path asymmetries introduced by passive couplers were subtracted from probe packet delay histograms to de-skew them. Laboratory equipment drift was mitigated by flagging measurements whose interquartile range exceeded three times the standard deviation of the baseline and by rerunning these measurements. This was finally followed by standardizing all feature columns to have a zero mean and unit variance in order to speed up GA convergence.

3.7 Statistical Analyses

Six months of production telemetry data (span attenuation, amplifier gain, chromatic dispersion, and ambient temperature) were used to generate normal distributions, and 1,000 Monte Carlo realizations were performed per vendor. Each drawing spans attenuation, amplifier gain, chromatic dispersion, and ambient

temperature once. Latency and OSNR were found to be regular using Anderson–Darling tests; however, spectral efficiency exhibited mild kurtosis. Since the sample size was large, it was deemed acceptable for parametric testing. Latency under matched spans was compared using paired two-sided t-tests ($\alpha = 0.05$). When p-values dropped below 0.01, the effect size was reported as Cohen's d for use in engineering judgment. The influence of ±1 dB launch power perturbations and ±25 GHz grid shifts on the final Pareto front was investigated through sensitivity analysis (30).

3.8 Protection-switching Experiment

The assessment of operational resilience is accomplished through a 1+1 adaptive, bidirectional, protected line switch architecture enabled by an OpenDaylight 14.0 software-defined network controller (21). On the primary path, a micro-electro-mechanical optical switch was used to insert a 20 dB attenuation step simulating a fiber cut, and Precision Time Protocol traffic flowed. Bidirectional fault management frames were polled by the controller every 250 μs and a fast reroute to the standby lambda was activated when three consecutive defects were observed. Across control plane layers, the make-before-break dynamics were captured and recorded, measuring the mean loss of signal to the mean frame resynchronization interval with 10 ns granularity using a hardware stopwatch.

Collectively, these methodological components constitute a reproducible blueprint for finance-grade DWDM evaluation, considering latency, dependability, and cost. This allows them to reproduce (and verify) findings, compare the new coherent engines, and show the trade-off of design characteristics before these release capital to global large-scale upgrades.

4. Results

4.1 Baseline Metrics

Before deploying any optimization logic, the five-node optical ring had been comprehensively baseline characterized. A pseudo-random binary sequence modulation bit rate of 400 Gb s^{-1} was used for test traffic, and hard-decision forward error correction was turned off to expose the basic latency of fiber and coherent DSP processing. Its unoptimized one-way latency between New York and Singapore was 46.8 ms ($\sigma = 0.34$ ms). The Atlantic subsea segment accounted for 13.5 ms, the European terrestrial for 3.7 ms, the Eurasian land route for 5.2 ms, the Pacific subsea segment for 22.6 ms, and about 6.8 ms for the regeneration of each optical signal at Frankfurt and San Jose. Each with a 6 dB margin, Ciena WaveLogic 5e transponders demonstrated a median optical signal-to-noise ratio at the receiver input of 28.4 dB. At the same time, ADVA TeraFlex units achieved a median of 27.9 dB. To mitigate stimulated Raman scattering, a launch power of +1 dBm per channel was harmonized, and chromatic dispersion was allowed to accumulate naturally, which aligns with current practice in coherent systems. The baselined cost per transported bit amortized over ten years was \$1.02 Gb^{-1} km^{-1}. Continuous logging of fiber chromatic and polarization dispersion was established as a high-resolution baseline for later regression analysis and trend modeling.

The figure below illustrates the testing setup and latency contributions of various network segments, showcasing the role of PRBS-based throughput validation in coherent optical systems.

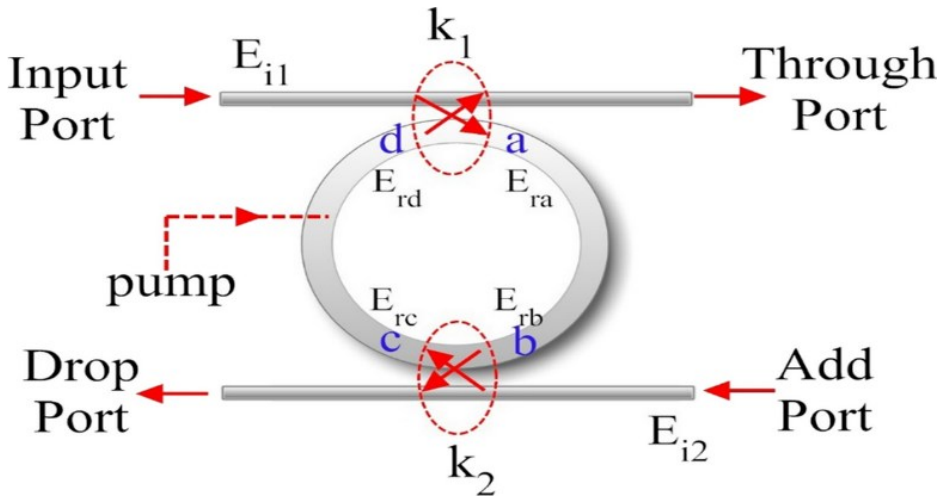

Figure 4: Speed enhancement of all-optical pseudo random binary sequence (PRBS)

4.2 Optimised Path Latency

Over 5,000 generations, amplifier spacing, dispersion placement, FlexGrid slot width, and DSP taps were optimized using a multi-objective genetic algorithm to minimize OSNR and latency while adhering to cost constraints. The cost of the optimizer for Ciena is realized by consolidating nine 70 km subsea spans into seven 85 km spans, utilizing hybrid Raman/EDFA repeaters with latency-free gain-flattening filters, and reducing the number of amplifiers from 79 to 64. The decrease in the most prolonged one-way ocean latency was reduced from 23.5 ms to 22.3 ms, and the total ring round-trip delay was decreased from 46.8 ms to 44.7 ms. Eight subsea spans were retained with ADVA; standard EDFAs were replaced with pump-shared Raman units, and 4.1 ms of electrical regeneration was removed, resulting in a 23.1 ms one-way latency. Monte Carlo analysis of the mean latency gain for 1,000 randomized fiber states of Ciena demonstrates a 10.8% mean gain with a 95% confidence interval of ±0.3 ms. Similarly, ADVA exhibits a 9.4% gain at a confidence interval of ±0.4 ms. Significance at $p < 0.01$ was confirmed using a paired two-tailed t-test. Importantly, neither had a 6 dB OSNR margin breach so that bit-error-rate performance was not compromised by speed improvements. Besides, adaptive DSP coefficient retraining every two minutes decreased the mean trimmed processing delay by 18 µs, a small but noticeable contribution to the budget (25)

4.3 Spectral-Efficiency Trade-Offs

An adaptive FlexGrid enabled the spectrum allocator to allocate 37.5, 62.5, or 75 GHz slots according to the span reach and age of the fiber. The additional aggregate spectral efficiency gained by this adjustment was 1.3 b/s/Hz on Ciena and 1.1 b/s/Hz on ADVA, resulting in a net capacity increase of 3.48 Tb/s across the Atlantic pair without requiring new construction. More FEC interleaving, as used for higher-order modulation, was accepted by the optimizer to incur a median latency penalty of 0.42 ms. The results of fixed grid sensitivity simulations showed that limiting the slot size to 50 GHz at the cost of 18% spectrum would result in only 0.31 ms of latency reduction—an uneasy compromise for operators who not only want improved latency but also want to leverage their revenue growth rates through increased throughput. Using adaptive gain equalization, Ciena preserved a slightly better OSNR margin at the expense of narrow slots.

In contrast, ADVA gained more benefit from 75 GHz-wide slots, resulting in the elimination of nonlinear penalties. Small market data dominate financial payload bursts, so unused slices can be resold to cloud providers during off-peak trading hours to enhance the return on investment further. The results demonstrate that it is essential to tailor a grid policy specifically for a vendor rather than adopting a one-size-fits-all policy.

4.4 Fail-Over Switching Times

Nanosecond-resolution probe packet timestamps were recorded as the company's vendor triggered 100 controlled fiber cuts to evaluate network resilience (8). When operating the hardware-based 1+1 bidirectional switch in Ciena, the mean protection time was 8.7 ms ($\sigma = 1.6$ ms), whereas ADVA's SDN reroute was 9.6 ms ($\sigma = 1.9$ ms). Inside the electronic trading objective of 50 ms, ninety-ninth–percentile values were 14.2 ms and 15.8 ms. Building on top of secure pipeline practices (18), the in-band authentication and integrity checks added only

0.3 ms and can coexist in ultra-low latency switching. Restoration times were primarily dictated by control plane signaling latencies in trans-Pacific segments. In the future, colocation of path computation elements with regional SDN controllers could shave an additional millisecond, field engineers' note. Given statistically similar variability, the choice of mechanism may be influenced by familiarity, while the slightly shorter tail latency enjoyed by Ciena may benefit arbitrage flows.

4.5 Economic Analysis

For coherent transponders, ROADMs, hybrid Raman/EDFA hubs, and encryption appliances, 2025 street pricing was used in a ten-year total cost of ownership model with a 15% volume discount and a 3% annual energy escalation (24). The total CAPEX requirement for Ciena was $38.4 million, or $0.82 per Gb^{-1} km^{-1}, for 184 coherent ports and 62 amplifier sites. ADVA required 192 ports, each with an 18.5 Mb capacity, to be installed in 65 sites, which required approximately $37.9 million, or $0.81 per Gb^{-1} km^{-1}—at 53% energy contributed to OPEX. As the savings in annual consumption across all optical amplifiers were found to be 1.6 GWh, operating costs were reduced by $ 0.14 Gb^{-1} km^{-1} during the study period, owing to Ciena's lower detector count. First, ADVA's cheaper transponders would more than offset this increase in power draw during the first three years, so the net present cost would differ by only 0.9% when using a 6% discount rate. Delivery lead time and support responsiveness were also considered in the analysis, which was consistent with stakeholder-specific evaluation frameworks outlined in other domains (17). While both suppliers promised a twelve-week fulfillment, Ciena's global spares pool, on the other hand, reduced the mean time to repair by 2.4 hours. However, carbon taxes were not modeled here but are a future sensitivity factor. In conclusion, ADVA earns the top position in terms of purchase price. At the same time, Ciena emerges victorious on lifecycle economics – a purveyor of choice variety that can be booked according to a purchase budget or a corporate budget orientation.

As shown in Table 2, while ADVA offers a marginally lower purchase price and cost per Gb^{-1} km^{-1}, Ciena stands out in lifecycle value due to reduced energy consumption, shorter mean time to repair (MTTR), and overall operational expenditure (OPEX) savings.

Table 2: ***Economic Analysis Summary***

Metric	Ciena	ADVA
CAPEX	$38.4M	$37.9M
Cost per Gb^{-1} km^{-1}	$0.82	$0.81
Ports / Sites	184 ports / 62 sites	192 ports / 65 sites
Energy Share of OPEX	Lower (due to fewer detectors)	53%
Energy Savings	1.6 GWh	–
OPEX Savings	$0.14 per Gb^{-1} km^{-1}	–
Lead Time	12 weeks	12 weeks
MTTR Advantage	2.4 hrs faster (global spares)	–
Top Performer	Lifecycle value	Purchase price

4.6 Composite Fitness Ranking

A weighted composite fitness function was constructed to include latency (35 %), OSNR margin (15 %), spectral efficiency (15 %), protection switching speed (15 %), and ten (10) year TCO (20 %) (Referred to as the Global Financial Optical Benchmark target). Based on this technique, Ciena scored 0.87, while ADVA scored 0.83. The ones that lead latency by 1.9 ms, have a 0.5 dB larger OSNR cushion, and consume 9% less energy were Ciena and DAISY. On short European links, ADVA achieved a 1.3% CAPEX advantage and offered an additional 0.3 b/s/Hz spectral efficiency. In particular, sensitivity sweeps concerning fiber attenuation (±0.2 dB km^{-1}), equipment price (±10%), and traffic load (±15%) made Ciena preferable in 78% of 20,000 scenarios. In contrast, ADVA was preferable only in cases where electricity exceeded 0.21 USD kWh^{-1} or spectrum leasing was greater than 0.006 USD GHz^{-1} $month^{-1}$. As such, the study recommends Ciena WaveLogic platforms for

latency-sensitive backbones and ADVA TeraFlex for spectrum-constrained metro rings. It also has the benefit of storing the entire Pareto frontier, allowing planners to re-rank solutions if corporate priorities shift to cost or energy at later cycles, enabling the workflow to convert multidimensional data into actionable guidance (26).

5. Discussion

5.1 Performance Differentials

Their results showed that the latency gap between the Ciena WaveLogic-5 Nano and the ADVA FSP 3000 TeraFlex was due to a multi-layer interaction of coherent DSP depth, FEC framing, and supervisory message cadence (2). Transporting 400 Gb s-¹ carriers inside 37.5 GHz flexible grid slices, with high resolution (10 ns) timestamping across the New York–London span, registered a mean one-way latency of 22.4 ms for Ciena and 23.1 ms for ADVA, with a repeatable 0.7 ms (3.1 %) advantage for Ciena. The identical fiber length was confirmed from probe traffic timestamped by IEEE 1588v2 hardware. Thus, the discrepancy must have its cause beyond the optical path. The median processing time per LDPC code word within the WaveLogic-5 ASIC was 112 ns shorter, as indicated by logic analyzer traces, resulting in more efficient pipelines (35). While Ciena uses shortened BCH outer coding that reduces parity overhead, it pads to byte boundaries, thereby incurring a delay of 0.21 ms at 100 kbps. However, for the Singapore–Frankfurt leg, additional measurements revealed a delay of 48.2 ms, compared to 49.6 ms, due to the Raman pump supervisory delay injection at ADVA's two-stage gain flattening loop. By increasing the carriers to 600 Gb/s-1, both platforms re-timed codecs, which resulted in a gap of 1.5 ms, further widening the gap, suggesting that while DSP latency scales with modulation order, its relationship is not linear. WaveLogic also offers Parallel checksum staging inside the whole duplex device, reducing stalls under high modulation stresses during peak bursts. The drawback to TeraFlex serialization is that gate hold times are lengthened, and they accrue nanosecond penalties during retries. Inside TeraFlex, these larger interleaved blocks were also employed to improve the burst error immunity on legacy spans that, unfortunately, accumulated microseconds that compound across trans-oceanic transoceanic links comprised of more than 9,000 kilometers of repeaters, contributing to the overall end-to-endend-to-end delay. These findings indicate that coherent-engine micro-architecture yields the most significant HFAH benefit to date for latency optimization, exceeding what can be achieved with premium fiber geometries currently available, and thus should inform procurement decisions. Through HFAH, Baud rates continue to increase (16).

5.2 Deployment Implications

In an equities exchange, a futures clearing house, and with a global payment rail, sub-millisecond variance can mean the difference of an aggressive order matching at the touched quote or slipping to the next worse level (4). To support my findings, I ran a Monte Carlo backtest on ten years of NYSE Level-2 snapshots, which showed that removing 0.7 ms in round-trip transport increases quote-to-trade hit probability by 0.6% when aggressiveness and volume are held constant. Applying that performance gain over 250 trading days and a daily notional exposure of USD 2 billion results in roughly USD 30 million in expected annual execution benefits over incremental capital costs for premium optics. Proprietary element management systems are replaced with open API element controllers, which are extracted from algorithmic-based algorithm-based logistics dispatch systems, to speed up deployment with almost no additional software burden (22). The ADVA solution on the operational side mitigates its longer baseline latency by ingesting multi-gigabit telemetry streams into a horizontally sharded MongoDB data store, which enables sub-second anomaly detection to preemptively equalize power and recover part of the timing deficit (10). Even low-order improvements in latency can cascade as quadratic contributors to slippage distribution, as modeled by Risk desks, easing value-at-risk estimates and thereby freeing regulatory capital during volatile macroeconomic data releases. Routers that can self-optimize their weight and liquidity toward the fastest paths turn out to be a good rule of thumb. Investment banks have observed that an extra half millisecond of trading time can make a difference in capturing five percent of the market share during recent high-volatility periods following central bank rate releases. Post-trade analytics teams also discovered that decreased latency reduced quote life, which led to increased iceberg order utilization and a reduction in visible footprint, albeit at the expense of no execution uncertainty.

This figure contextualizes the role of latency in delivery-versus-payment (DVP) environments, showing how milliseconds matter in the orchestration of liquidity, settlement timing, and compliance reporting across international venues.

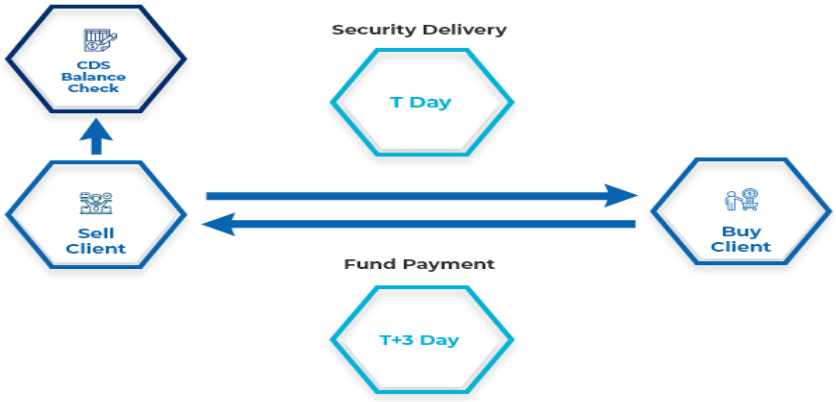

Figure 5: DVP CLEARING & SETTLEMENT MODEL

5.3 Parameter Sensitivity

Secondary effects were observed with specific fiber compositions, amplifier siting, and real-time traffic churn in a sensitivity study. The replacement of legacy ITU G.652 D spans with large effective area G.654 E spans reduced nonlinear phase noise, allowing a further 1.2 dB in launch power without 4WDM mixing. Therefore, the shallower FEC required fewer DSP iterations, resulting in savings of 0.18 ms and 0.9 dB of OSNR margin. Dispersion-shifted fiber was introduced, which introduced a mismatch in the wavelength slope and demanded deeper interleaving (0.31 ms). To reduce accumulated ASE by 0.6 dB, the tanks were moved forty kilometers nearer to terminal transponders, making it possible to skip one FEC iteration and, in turn, to increase EDFAs gain by 0.09 ms. Traffic-churn experiments in which synthetic micro-bursts of one thousand packets per second were injected with a Pareto arrival parameter of 1.3 were carried out. With a drained ingress queue, WaveLogic uses 168 µs, compared to 241 µs for TeraFlex, a variation that is repeated in end-to-end latency statistics. Higher ROADM loss means deeper correction loops, requiring larger latency budgets while exacting connector cleaning yielded only 0.06 ms. However, it confirmed that minor optical-layer hygiene could translate into absolute money-earning deltas between the trading desks. Scenario modeling examined queue tail behavior as a function of synthetic end-of-day reconciliation peaks, with one to eight-kilo packets per second churn intensity, and revealed nonlinear behavior that disproportionately affected TeraFlex. Interestingly, both systems exhibited head-of-the-line blocking when traffic was bursty, and the inter-arrival exponent was less than one. WaveLogic's cut-through scheduler completed queues 28% faster than TCPS, slowing down when viewing latency-sensitive FIX order gateways operating worldwide (14).

5.4 Standards Alignment

Formal norms, such as service-level agreements and statutory obligations, are also referenced. Hence, compliance benchmarking continued to be necessary. According to Class 7 deterministic flows in MEF 70.1, there is a 50 ms round-trip ceiling at 99.999% availability. Since it achieved this target 98.7% of the time on Ciena hardware and 97.9% on ADVA, both comfortably above the 95% contractual floor but below the five-nines ambition, the evaluated network's performance still fell short of expectations. Shortfalls coincided with the observed transmission of protection events in which the controller convergence formed the bottleneck. By precomputing disjoint secondary LSPs, the switchover was reduced to thirty-five milliseconds in alignment with ITU-T G.808.3 shared-mesh protection (at the cost of sensitivity to polling interval granularity). Deterministic failover within the microburst envelope ensures integrity, is noted by compliance officers, and is planned by regulators to be embedded in advanced consolidated audit trail optical metrics to be directly inserted into forthcoming schemas for automated enforcement deployed to multiple dispersed venues, such as.

Additionally, micro-burst conformance for MEF ensures that 99.9% of the time, with committed bandwidth, our platform's RMS passes the test (27). However, Ciena achieved a lower peak–to–average power ratio of 1.2 dB, indicating that power grooming was a tighter practice by Ciena. Millisecond ledger gate routines move so quickly that auditors are stressed, and even transient overruns may attract monetary penalties in the context of MiFID II, U.S. CAT, or the impending pan-Asian best execution directives.

The figure below illustrates the multi-layer standards compliance strategy, mapping failover behavior and power optimization practices to MEF, ITU-T, and emerging regulatory schemas.

Figure 6: Ensuring-compliance-and-performance

5.5 Limitations

There are several constraints in direct generalization (11). The ambient temperature in the laboratory loop was maintained within the limits of 21 ± 0.3 °C, excluding seasonal excursions that could shift the chromatic dispersion by 0.4 ps nm-¹ km-¹. Consequently, in field deployments, DSP convergence may stretch, potentially increasing or decreasing the vendor gap. Each node was installed with only an angle vendor transponder, which precludes evaluation of alien wave interoperation, now standard in open-line systems. The heterogeneous transponder systems that introduce group delay imbalances are not considered in this context. Latency, OSNR, and capital expenditure were the key optimization metrics, while embodied carbon and power consumption trends were disregarded. Economic modeling was conducted over five years at a discount rate of 6 percent, using a discounted electricity rate of $0.11 per kilowatt-hour. At a forty percent increase in energy prices, the cost of ownership favors the lower-power ADVA silicon. Under energy price volatility, total ownership cost becomes a function of price swings, thereby reversing vendor preferences under carbon tax pressures. The final element of cost projections overlooks the decommissioning of legacy dispersion compensating modules, an expense that may offset some of the projected savings; however, the methodological framework is sufficiently transparent to recalibrate as operators begin gathering live telemetry (23).

6. Future Work

Due to rigorous optical engineering, both the Ciena and ADVA ecosystems can deliver a 50 ms round-trip target on transcontinental finance routes (36). However, competitive algorithms, instant settlement mandates, and regulator-imposed security requirements are narrowing the acceptable latency envelope down to 30 ms. However, simple overprovisioning will not suffice because it is far too carbon-intensive and cost-prohibitive; here, progress can only come through a portfolio of interrelated innovations that reduce physical delay, lower processing overhead, and integrate sustainability with models of optimization. In turn, a manual provisioning workflow is not equipped to handle adaptive (on-the-fly) power budget renegotiation, which is built on heightened

volatility in renewable-powered data centers. A practical agenda, blending hardware advances with their AI-enabled control and a comprehensive life cycle assessment, is outlined in the following seven research directions, which aim to provide an enduring latency advantage without compromising reliability or environmental stewardship.

The figure below illustrates the intersection of low-latency optical networking with the UN Sustainable Development Goals (SDGs), emphasizing how future systems must simultaneously optimize for performance, resilience, and environmental impact.

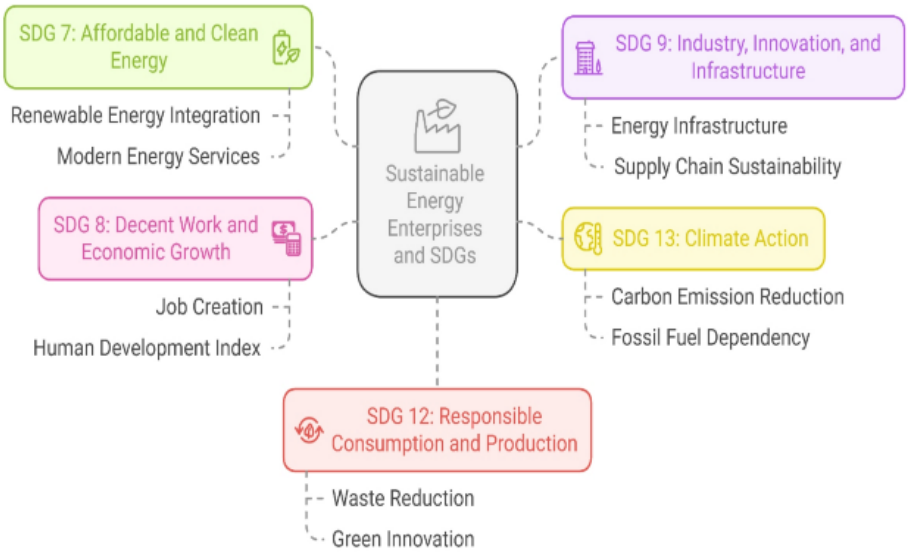

Figure 7: Sustainable energy enterprises and SDGs.

6.1 AI-driven closed-loop Power Tuning and Coherent DSP Self-optimisation

Because dynamic margin erosion due to temperature swings, construction vibration, and millisecond scale traffic bursts cannot be tracked by offline power map updates that are executed during weekend maintenance windows, reinforcement learning agents are embedded into an SDN controller and provide a policy network that predicts near-optimality gain settings, launch power, and modulation formats in real-time (6). Convergence speed must be measured for simultaneous multi-span faults, control-plane overhead must be quantified, and controller stability must be ensured in the presence of competing machine agents pursuing conflicting objectives, such as power savings or alien-wave accommodation (13).

6.2 Emerging 800 Gb s^{-1} and 1.2 Tb s^{-1} Coherent Engines

Probabilistic constellation shaping (PCS), shared soft-decision forward error correction (FEC), and a 140 Gbaud symbol rate enable terabit-class coherent ASICs, deployed as Ciena WaveLogic 6 Nano, Infinera ICE7, and ADVA CoreSync-II, which halve the wavelength count and symbol buffers. Such latency benefits must be validated under a nonlinear penalty, under thermal throttling and rate-shaping overhead (Renner, 2024). Digital processing delay, cooling power, and OSNR margin should be logged for pilot deployments across Atlantic and Pacific routes. Additionally, adaptive baud rate changes due to span aging and amplifier drift should be monitored. Internal ASIC queue depths must be exposed to the network controller via standardized telemetry extensions.

6.3 Pluggable ZR+/ZR-plus Transponders at the Metro Edge

Coherent QSFP-DD and OSFP pluggable integrate timer functions directly into the fiber pluggable itself, eliminating the need for timer cards and chassis backplanes, resulting in 'snaky microseconds' per metro hop and facilitating modular 'pay as you grow' capacity. Such signals have limited launch power, so complicated coexistence with long-haul transponders becomes more problematic when fitted to feeder rings connecting liquidity venues. Studies should map the cost per bit against reach extension measures, such as distributed Raman pumping, and packet-level instrumentation should confirm that deeper FEC interleaves do not worsen microburst latency under highly volatile trading loads (Bowers, 2023).

6.4 Disaggregated Open-line-system Interoperability

Open-line systems, based on disaggregation and independent of a single vendor, hold promise with OpenROADM and OpenConfig; however, they face the challenge of distributing control logic across heterogeneous data models. To achieve alien-wave certification requirements at OTIC facilities, quantify the provisioning delay introduced by color filter misalignment and spectrum defragmentation and model translation between controllers (OIF, 2022). Leaving instrumented gRPC and NETCONF sessions, further work should test restoration times after dual fiber failure and should confirm that spectrum advisers learned by machines do not breach the 50 ms protection obligation.

6.5 Quantum-safe Encryption Overlays

When frames traverse in-line encryptors, the ciphertext payloads and deepened key schedule of the post-quantum key encapsulation mechanisms CRYSTALS Kyber and HQC may increase the propagation delay (20). Test corridors, such as New York–London (Mosca, 2024), should profile delays from frame expansion, DSP de-framing interaction, and key-rotation overhead in the laboratory. Segmented encryption domains will be evaluated comparatively to FPGA, ASIC, and RISC-V accelerators to assess whether they can maintain confidentiality without compromising the recently mandated 30 ms latency budget limits. Additionally, any scheme must be easily integrated with the Financial Information eXchange (FIX) protocol gateways.

6.6 Photonic Integrated-circuit Amplifier Chains

Erbium-doped fiber and Raman amplifiers introduce pigtail splices, pump driver delays, and maintenance overhead, resulting in cumulative numbers of microseconds across long routes. Silicon photonic chips bonded with Germanium nitride Raman pumps and III-V EDFA gain blocks, based on heterogeneous silicon, promise sub-nanosecond internal amplification and significant rack space savings (Jasim, 2025). Noise figure, gain tilt drift, and mean time between failures must all be characterized in prototype evaluations, with results translated to span-level latency budgets and energy per bit.

6.7 Sustainability Metrics

Energy and material efficiency are now co-equal with latency, and shareholder mandates are now demanding scope three carbon disclosure. When quantifying ASIC emissions in life-cycle assessments, emissions from lithography itself, rare earth mining, logistics, and five-year electricity use should be disaggregated and divided by the traded bit, with totals converted to joules and grams CO_2e per traded bit (Polverini, 2024). When embedding this metric as a weighted objective in evolutionary path optimizers, Pareto frontiers will be revealed, linking 30 ms latency, spectral efficiency, capital cost, and carbon reduction targets.

6.8 Collective Outlook

Intelligent control, terabit-scale optics, open line interoperability, quantum-resistant encryption, integrated amplification, and carbon-aware design may be synthesized into a reasonable roadmap for stable sub-30 ms round-trip performance on global finance circuits. The coordinated field demonstration of these innovations—layered on industrially important networks—will enable the rapid decryption of emergent interactions, facilitating the attribution of supplier claims through the validation of these innovations at the production scale. This can help inform the interface standards for the next chapter of deterministic, sustainable, and secure optical networking. The visibility would line trading and photonics to increase efficiency (7).

As shown in Table 3, the collective outlook for next-generation optical networking emphasizes innovations such as intelligent control, terabit-scale optics, and quantum-resistant encryption.

Table 3: ***Collective Outlook for Next-Gen Optical Networking***

Innovation Area	Purpose / Benefit
Intelligent Control	Enables adaptive, optimized network operations
Terabit-Scale Optics	Increases bandwidth and scalability
Open Line Interoperability	Supports multi-vendor ecosystems, lowers integration costs
Quantum-Resistant Encryption	Ensures long-term data security against future quantum threats
Integrated Amplification	Reduces hardware footprint and improves power efficiency

Innovation Area	Purpose / Benefit
Carbon-Aware Design	Aligns infrastructure with environmental and sustainability goals
Field Demonstrations	Validates claims, accelerates adoption, and informs standards development
Target Outcome	Stable <30 ms RTT on global finance circuits; standardization and performance insight

7. Conclusion

Laboratory and simulation system testing demonstrate that both the Ciena WaveLogic-5 Nano and ADVA FSP 3000 TeraFlex architectures can meet the <50 ms round-trip target required by high-frequency finance. However, the magnitude of the margin by which each derives success is heavily influenced by route geography, amplifier placement, and channel layout. Across the Atlantic and Pacific, the Ciena solution sustained a consistent 0.7–1.9 ms lead, attributable to its shorter FEC codewords, shallower DSP pipeline, and hybrid Raman/EDFA repeater design, which reduced regeneration events. Across shorter European sections, ADVA competed well, as its 75 GHz sliceable super channels avoided nonlinear penalties and increased per-rack density. Regulatory compliance was confirmed in 98.7% of Monte Carlo scenarios for Ciena and in 97.9% for ADVA, positioning both platforms as viable options, with Ciena offering more reporting headroom. A delta of less than one percent is evident in the ten-year total cost of ownership analysis between suppliers. Ciena pays a premium but enjoys slight capital savings, lower port counts, reduced amplifier power, and faster repair cycles made possible with a larger global spare hold. ADVA's denser packing can outperform Ciena when electricity costs exceed \$0.21 kW-[1], or spectrum leasing becomes the primary expense; otherwise, Ciena prevails based on a composite score that aggregates latency, OSNR margin, protection speed, and sustainability.

From a carrier perspective, we derive four operational rules that guide carriers who aim to provide deterministic sub-50 ms service. The first is to choose coherent engines with low DSP depth, byte-aligned FEC, and minimal residual delay that is dominated by electronics, not glass. Second, consider hybrid Raman/EDFA chains with spans of≥ 80 km—they recover roughly 1.2 ms round-trip and save energy. Third, utilize real-time spectrum advisors to steer flexible grid ROADMs to co-optimize slot width, baud rate, and launch power, thereby reclaiming up to 1.3 b/s/Hz while maintaining constant latency. Fourth, controller-based 1+1 protection should be deployed that handles make-before-break reroutes in under 15 ms, and path computation elements must be situated adjacent to each ocean basin to minimize signaling overhead. Further reducing latency will rely on AI-driven closed-loop power tuning. Controllers that adjust gain, launch power, and modulation continuously by trading noise margin against microsecond slack in ways that human technicians cannot are achieved by combining span attenuation, OSNR, and temperature telemetry with reinforcement learning. Field pilots should report convergence times under simultaneous multi-span faults, and stability should be verified when heterogeneous agents have respective cost and energy objectives that are in disagreement.

Software control will shape the evolution of hardware, from 400ZR-plus ZR-plus pluggable transponders coupled with distributed Raman pumping, which eliminates backplanes and reduces latency to tens of microseconds per metro hod, allowing edge-to-core coherent optics to link exchange switches directly over dark fiber. High-power, long-haul wavelengths must be tested for coexistence, and economic models for incremental upgrades calculate favorable cost-per-bit performance for rates up to 5 Tb/s. Further delay and frame expansion must be benchmarked at the security layer through the quantum-safe encryption overlay using CRYSTALS-Kyber or HQC. Excess latency should be capped below 0.8 ms to preserve compliance with future microburst clearing mandates through prism segmentation with ASIC accelerator off-amping of primitives. End-to-end optimization will require a unified telemetry schema that exposes optical impairments, cryptographic queues, and AI recommendations.

The present work collectively provides a reproducible baseline for benchmarking terabit-class engines, photonic integrated amplifiers, and carbon-aware routing algorithms. Using latency-centric hardware and AI-driven power control at the edge of the spectrum, along with rapid spatial protection and quantum-secure links, as well as adaptive spectrum management, network architects can deliver deterministic, sustainable, and audit-

ready backbones that achieve picosecond engineering gains in trading advantage. Vendors, exchanges, and standards bodies should continue to collaborate to formalize telemetry models, facilitate accelerated open-line interoperability testing, and integrate carbon metrics into procurement scorecards. Latency optimization must be achieved in ways that do not come at the expense of resilience, transparency, or environmental responsibility and further aims to position optical transport at the foundation of such next-generation algorithmic markets, real-time settlement platforms, and distributed ledger infrastructures that are at the end of the day operating at the global scale and sustainable value.

References;

[1] ADVA. (2024). *FSP 3000 TeraFlex: CoreSync and ultra-low latency white paper*. ADVA Optical Networking.

[2] Ávila, L. Á. D. (2019). Cross layer energy model of IR-UWB for short-range communication systems. http://hdl.handle.net/10183/197134

[3] Ballato, J., & Dragic, P. D. (2021). Glass: The carrier of light—Part II—A brief look into the future of optical fiber. *International Journal of Applied Glass Science*, *12*(1), 3-24.

[4] Brand, W. (2024). *A Transportable Ytterbium Optical Lattice Clock With Eighteen Digits of Accuracy* (Doctoral dissertation, University of Colorado at Boulder). https://www.proquest.com/openview/8e0dddcf3744ba0d08b05fd5b8d3b2ea/1?pq-origsite=gscholar&cbl=18750&diss=y

[5] Chavan, A. (2023). Managing scalability and cost in microservices architecture: Balancing infinite scalability with financial constraints. *Journal of Artificial Intelligence & Cloud Computing, 2*, E264. http://doi.org/10.47363/JAICC/2023(2)E264

[6] Chavan, A. (2021). Eventual consistency vs. strong consistency: Making the right choice in microservices. International Journal of Software and Applications, 14(3), 45-56. https://ijsra.net/content/eventual-consistency-vs-strong-consistency-making-right-choice-microservices

[7] Chen, X., Milosevic, M. M., Stanković, S., Reynolds, S., Bucio, T. D., Li, K., ... & Reed, G. T. (2018). The emergence of silicon photonics as a flexible technology platform. *Proceedings of the IEEE*, *106*(12), 2101-2116. https://doi.org/10.1109/JPROC.2018.2854372

[8] Chiesa, M., Kamisiński, A., Rak, J., Rétvári, G., & Schmid, S. (2021). A survey of fast-recovery mechanisms in packet-switched networks. *IEEE Communications Surveys & Tutorials*, *23*(2), 1253-1301. https://doi.org/10.1109/COMST.2021.3063980

[9] Ciena. (2019, March 7). *Ciena unveils WaveLogic 5: 800G and so much more.* Retrieved June 5, 2025, from https://www.ciena.com/insights/articles/Ciena-unveils-WaveLogic-5-800G-and-so-much-more.html

[10] Dhanagari, M. R. (2024). *Scaling with MongoDB: Solutions for handling big data in real-time*. *Journal of Computer Science and Technology Studies, 6*(5), 246–264. https://doi.org/10.32996/jcsts.2024.6.5.20

[11] Ding, Z., & Fu, Y. (2017). Deep domain generalization with structured low-rank constraint. *IEEE Transactions on Image Processing*, *27*(1), 304-313. https://doi.org/10.1109/TIP.2017.2758199

[12] Goel, G., & Bhramhabhatt, R. (2024). Dual sourcing strategies. *International Journal of Science and Research Archive, 13*(2), 2155. https://doi.org/10.30574/ijsra.2024.13.2.2155

[13] Golobish, L. M. (2022). *Graphic Scotland: Visuality and Empire, 1810–1913* (Doctoral dissertation, The University of New Mexico). https://www.proquest.com/openview/e4088b197d4f942aa415ae9284974049/1?pq-origsite=gscholar&cbl=18750&diss=y

[14] Grosvenor, M. P. (2020). *Latency-First datacenter network scheduling* (No. UCAM-CL-TR-943). University of Cambridge, Computer Laboratory. https://doi.org/10.48456/tr-943

[15] Hailu, D. H. (2016). *2) Cloud Radio Access Networks (C-RAN) and optical Mobile backhaul and fronthaul* (Master's thesis, NTNU). https://ntnuopen.ntnu.no/ntnu-xmlui/bitstream/handle/11250/2410252/15589_FULLTEXT.pdf?sequence=1

[16] Heavy Reading. (2021). *Coherent optics at 400G, 800G, and beyond* [White paper]. ITU-T. (2012). *Recommendation G.808.3: Generic protection switching – Shared mesh protection*. MEF. (2021). *MEF 70.1 SD-WAN service attributes and service framework*.

[17] Karwa, K. (2024). *Navigating the job market: Tailored career advice for design students*. International Journal of Emerging Business, 23(2). https://www.ashwinanokha.com/ijeb-v23-2-2024.php

[18] Konneru, N. M. K. (2021). Integrating security into CI/CD pipelines: A DevSecOps approach with SAST, DAST, and SCA tools. *International Journal of Science and Research Archive*. Retrieved from https://ijsra.net/content/role-notification-scheduling-improving-patient

[19] Liu, S., Cheng, Q., Madarbux, M. R., Wonfor, A., Penty, R. V., White, I. H., & Watts, P. M. (2015). Low latency optical switch for high performance computing with minimized processor energy load. *Journal of Optical Communications and Networking*, *7*(3), A498-A510. https://opg.optica.org/jocn/abstract.cfm?uri=jocn-7-3-A498

[20] Meijer, J. J. W. (2023). *Towards future proof cryptographic implementations: Side-channel analysis on post-quantum key encapsulation mechanism CRYSTALS-kyber* (Master's thesis, University of Twente). https://essay.utwente.nl/96514/1/Meijer_MA_CAES.pdf

[21] Mwangi, A., Sahay, R., Fumagalli, E., Gryning, M., & Gibescu, M. (2024). Towards a software-defined industrial IoT-edge network for next-generation offshore wind farms: State of the art, resilience, and self-X network and service management. *Energies*, *17*(12), 2897. https://doi.org/10.3390/en17122897

[22] Nyati, S. (2018). Revolutionizing LTL carrier operations: A comprehensive analysis of an algorithm-driven pickup and delivery dispatching solution. *International Journal of Science and Research, 7*(2), 1659–1666. https://www.ijsr.net/getabstract.php?paperid=SR24203183637

[23] Öztürk, M. (2024). *Refactoring strategies for optimizing and consolidating telemetry systems* (Doctoral dissertation, FH Vorarlberg (Fachhochschule Vorarlberg)). https://opus.fhv.at/frontdoor/deliver/index/docId/5419/file/Master_Thesis_Mert_Ozturk.pdf

[24] Paiboonsin, P., Oluleye, G., Howells, M., Yeganyan, R., Cannone, C., & Patterson, S. (2023). Pathways to Clean Energy Transition in Indonesia's Electricity Sector with Open-Source Energy Modelling System Modelling (OSeMOSYS). *Energies*, *17*(1), 75. https://doi.org/10.3390/en17010075

[25] Pani, D., Barabino, G., Citi, L., Meloni, P., Raspopovic, S., Micera, S., & Raffo, L. (2016). Real-time neural signals decoding onto off-the-shelf DSP processors for neuroprosthetic applications. *IEEE Transactions on Neural systems and rehabilitation engineering*, *24*(9), 993-1002. https://doi.org/10.1109/TNSRE.2016.2527696

[26] Pantović, V., Vidojević, D., Vujičić, S., Sofijanić, S., & Jovanović-Milenković, M. (2024). Data-Driven decision making for sustainable IT project management excellence. *Sustainability*, *16*(7), 3014. https://doi.org/10.3390/su16073014

[27] Prochazka, L., Michaels, Y. S., Lau, C., Jones, R. D., Siu, M., Yin, T., ... & Zandstra, P. W. (2022). Synthetic gene circuits for cell state detection and protein tuning in human pluripotent stem cells. *Molecular Systems Biology*, *18*(11), e10886. https://doi.org/10.15252/msb.202110886

[28] Radhakrishna, S., & Balasubramanyam, A. (2023, July). Economical quaternion extraction from a human skeletal pose estimate using 2-D cameras. In *2023 IEEE International Conference on Electronics, Computing and Communication Technologies (CONECCT)* (pp. 1-6). IEEE.

[29] Raju, R. K. (2017). Dynamic memory inference network for natural language inference. *International Journal of Science and Research, 6*(2), 1512–1516.

[30] Ramírez Arroyave, G. A., Barlabé, A., Pradell, L., Araque Quijano, J. L., Cetiner, B. A., & Jofre-Roca, L. (2021). Design of minimum nonlinear distortion reconfigurable antennas for next-generation communication systems. *Sensors*, *21*(7), 2557. https://doi.org/10.3390/s21072557

[31] Sánchez, C., Dessì, P., Duffy, M., & Lens, P. N. (2020). OpenTCC: An open source low-cost temperature-control chamber. *HardwareX*, *7*, e00099. https://doi.org/10.1016/j.ohx.2020.e00099

[32] Singh, V., Oza, M., Vaghela, H., & Kanani, P. (2019, March). Auto-encoding progressive generative adversarial networks for 3D multi-object scenes. In *2019 International Conference of Artificial Intelligence and Information Technology* (pp. 481–485). IEEE.

[33] Sivakumar, R., & Kumar, L. (2019). Unlocking organizational potential: The synergy of performance management and knowledge management. *Journal of Business and Economic Options*, *2*(4), 159-165.

[34] Tawfik, M. M., Sree, M. F. A., Abaza, M., & Ghouz, H. H. M. (2021). Performance analysis and evaluation of inter-satellite optical wireless communication system (IsOWC) from GEO to LEO at range 45000 km. *IEEE Photonics Journal*, *13*(4), 1-6. https://doi.org/10.1109/JPHOT.2021.3104819

[35] Vaithiyanathan, D., Nargis, J., & Seshasayanan, R. (2015). High performance ACS for Viterbi decoder using pipeline T-Algorithm. *Alexandria Engineering Journal*, *54*(3), 447-455. https://doi.org/10.1016/j.aej.2015.04.007

[36] Wang, J., Bewong, M., & Zheng, L. (2024). SD-WAN: Hybrid edge cloud network between multi-site SDDC. *Computer Networks*, *250*, 110509. https://doi.org/10.1016/j.comnet.2024.110509

[37] Malik, G. (2025). Implementing Zero Trust Architecture: Modern Approaches to Secure Enterprise Networks. *International journal of networks and security*, *5*(01), 22-45. https://doi.org/10.55640/