Delay-Aware Time-Coordinated Nanosecond Optical Switching for Dynamically Reconfigurable Optical Networks

Ali Mehrpooya*, Vaigai Nayaki Yokar*, Yiran Teng*, Sen Shen*, Wanxin Zhao*, Arash Olianezhad*, Sima Bahrani*, Amin Emami*, Ning Zhang†, Shuangyi Yan*, Dimitra Simeonidou*

*HPN Group, Smart Internet Lab, University of Bristol, Bristol, United Kingdom

†Compound Semiconductor Applications Catapult, United Kingdom

{shuangyi.yan}@bristol.ac.uk

***Abstract*—We demonstrate an FPGA-embedded AI control platform enabling real-time, nanosecond-accurate optical path reconfiguration. Using an on-chip MLP model for delay prediction, our system reduces packet loss by 90%, achieving deterministic, self-healing switching for future AI-native optical networks.**

***Index Terms*—Precision Time Protocol (PTP), Optical switching, Software Defined Networking (SDN), Network synchronization, High-speed switching, Time synchronization, FPGA**

I. Introduction

Future optical transport networks are evolving toward *self-managed* and *self-healing* architectures capable of reconfiguring links and topologies in real time to respond to traffic dynamics, failures, or anomalies [1], [2]. In these intelligent systems, fast optical switches play a central role in enabling dynamic topology changes. Recent work has primarily focused on time-synchronized optical switching, ensuring that distributed nodes switch simultaneously, ignoring the impact of fiber propagation delay between switching points. [3]–[5]. As a result, packets in transit during reconfiguration can still be lost, limiting the effectiveness of synchronization-only schemes and introducing short service interruptions during topology updates. To achieve lossless and deterministic reconfigurations, both precise time synchronization and accurate scheduling of switching events must operate together.

In this paper, we propose a delay-aware optical switching architecture that performs real-time delay estimation and intelligent scheduling based on precise time synchronization using the Precision Time Protocol (PTP). Regenerated clock signals derived from synchronized 1PPS timing enable switching operations with 4 ns resolution. To achieve lossless reconfiguration, an execution-time offset between two optical switches is dynamically set based on the measured propagation delay of the optical path. FPGA-based control agents continuously exchange PTP timestamps to measure path delay, while a lightweight Multi-Layer Perceptron (MLP) model maps these measurements to accurate downtime predictions. Both delay estimation and scheduling are implemented directly in FPGA hardware, ensuring sub-microsecond control latency and enabling autonomous switching decisions with nanosecond precision. By jointly optimizing synchronization and per-link scheduling, the proposed system reduces packet loss by up to 90% during optical reconfiguration of two node optical route for both 1 Gb/s and 10 Gb/s data rates. The FPGA-driven, AI-assisted switching scheme establishes a practical foundation for deterministic, self-healing optical networks aligned with future 6G and AI-native transport infrastructures [6].

II. Validation System and Experimental Setup

Figure 1 illustrates the principle of propagation-delay–aware optical switching. At time t_0, as shown in Fig. 1(a), packets p_1–p_n propagate along Path 1. If both switches simultaneously reconfigure to Path 2 at time t_1 [Fig. 1(b)], packets still propagating along Path 1 are lost, and the receiver on Path 2 temporarily receives no signal because the recently emitted packets have not yet arrived at the downstream switch. This synchronized switching, although aligned in time, leads to momentary data loss due to unaccounted fiber delay. As shown in Fig. 1(c-d), a lossless transition is potentially achieved by introducing an offset at the downstream switch, delaying its reconfiguration to $t_1 + \Delta t$, where Δt corresponds to the propagation delay. This ensures that the switch activates precisely when the first packet from the new path arrives, maintaining signal continuity and minimizing packet loss. To achieve this, each node must maintain a precisely synchronized local clock across the distributed system. Fast optical switches with nanosecond-level reconfiguration time must be coordinated according to their local time reference through FPGA-based control agents. The SDN controller oversees this coordination, issuing global switching intents and delegating accurate exe-

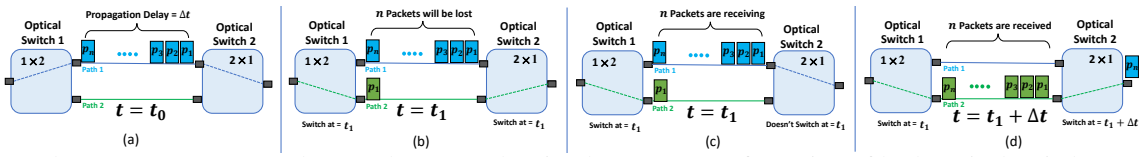

Fig. 1: (a) Packets p_1–p_n propagate along Path 1 at t_0. (b) Simultaneous reconfiguration of both optical switches at t_1 interrupts in-flight packets, causing loss. (c) The downstream node waits for the propagation delay Δt. (d) Switching at $t_1 + \Delta t$ aligns with packet arrival at Path 2, achieving lossless transition.

TABLE I: Summary of the trained MLP model layers parameters.

Layer	Shape	Act.	Params
Input	1	–	–
Hidden 1	8	ReLU	16
Hidden 2	8	ReLU	72
Output	1	Linear	9
Total			**97**

TABLE II: Cross-validation on mixed real and synthetic training data.

Dataset	MSE	MAE	R^2
Training	0.098	0.243	0.995
Validation	0.112	0.265	0.992
Testing	0.103	0.265	0.993

TABLE III: Averaged packet-loss comparison.

Rate	Condition	Loss	Reduction
1 Gb/s	Conventional	25.4	–
1 Gb/s	AI-Assist	2.4	90.6%
10 Gb/s	Conventional	226.2	–
10 Gb/s	AI-Assist	21.5	90.5%

cution timing to local agents to ensure deterministic, lossless reconfiguration.

The experimental setup is shown in Fig. 2. Two Raspberry Pi Compute Module 4 (CM4) units, each equipped with Open TimeCards, serve as PTP-synchronized time agents. These agents exchange PTP messages through a PTP-enabled Ethernet switch, while their optical path coincides with the same fiber infrastructure used for real data traffic. This design enables simultaneous clock synchronization between nodes and estimation of the end-to-end propagation delay along the actual data-carrying fiber. Each agent provides 1 PPS and 10 MHz reference signals, which are fed into a Xilinx VCU108 Virtex UltraScale FPGA. The FPGA derives a 250 MHz internal clock for 4 ns scheduling precision. Three modules are implemented in the FPGA: (i) a time-synchronizer and scheduler, (ii) a real-time MLP-based delay estimator, and (iii) an optical switch controller. The system supports two operation modes: a *manual mode*, where the user inputs the measured propagation delay, and an *automatic mode*, where delay is estimated online from PTP measurements. Since raw PTP delay measurements include additional latency from Ethernet switches, EDFAs, and SFP+ transceivers, they cannot directly represent the true fiber propagation delay with nanosecond accuracy. To address this, a lightweight Multi-Layer Perceptron (MLP) with two hidden layers (eight neurons each) is implemented on the FPGA to learn and compensate for these systematic delays. Trained offline using both synthetic and measured PTP datasets, the model maps incoming PTP timestamp data to accurate per-path delay estimates. Running natively in hardware, it enables dynamic adaptation to link conditions and nanosecond-precision scheduling of switching events. Two 1×4 magneto-optical optical switches are connected via four 50 km SMF links. The data plane is driven by an Anritsu Data Analyzer generating 1 Gb/s and 10 Gb/s IPv4 traffic (1000-byte and 1504-byte packets). Data and PTP signals are multiplexed, amplified by an EDFA, and transmitted through the shared fiber. A high-speed photodiode connected to the output of the second optical switch monitors optical power during reconfiguration. Any transient drop in photodiode voltage corresponds to missing packets or downtime during switching. By compensating for the measured propagation delay, the FPGA minimizes these downtime and packet loss, achieving near-continuous transmission during path reconfiguration. The SDN controller (PYNQ-Z1) supervises the overall experiment and calibrates both FPGA nodes for Ethernet delay alignment.

III. Experimental Results

As shown in Fig. 2, the proposed system operates in two modes: manual and AI-assisted. In manual mode, the user programs the FPGA scheduler to trigger switching at a specific time, providing flexibility for controlled experiments. In AI-assisted mode, a trained MLP model predicts the optical

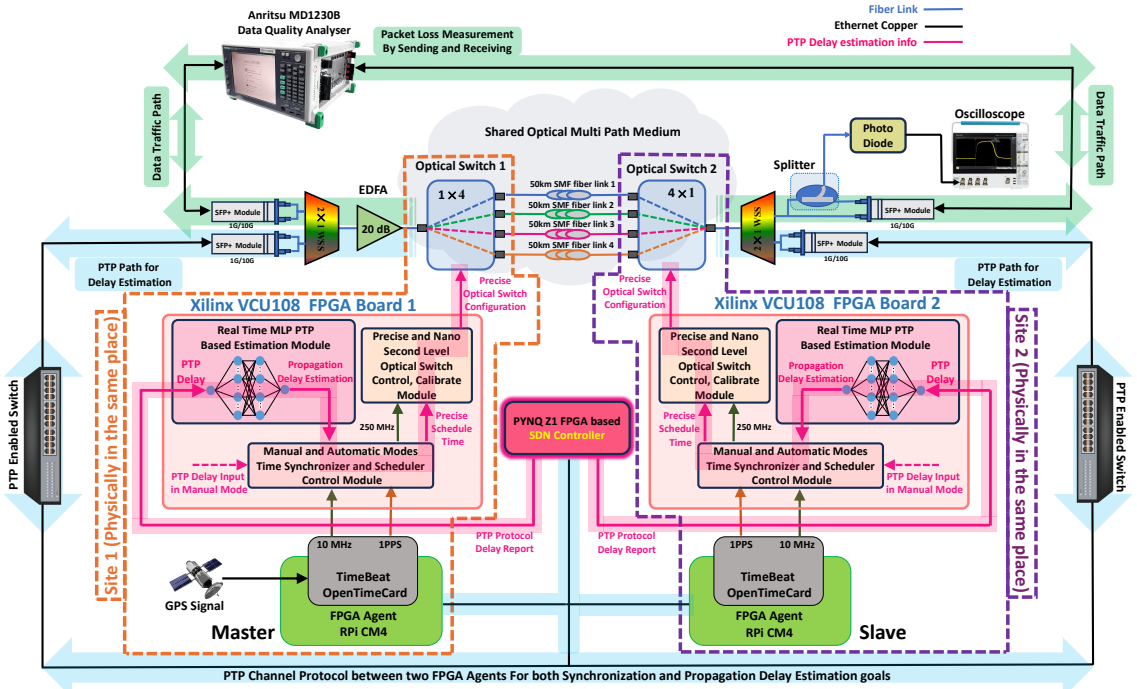

Fig. 2: Experimental setup for manual and AI-assisted modes, propagation-delay–aware optical switching.

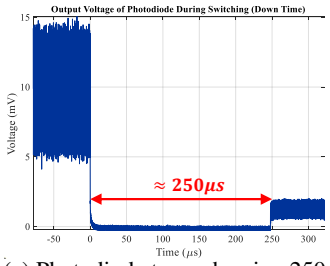

(a) Photodiode trace showing 250 µs downtime due to 50 km propagation delay.

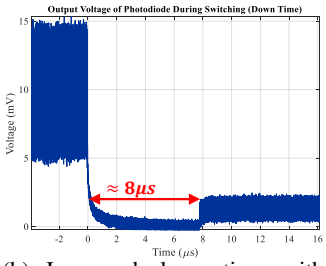

(b) Improved down time with 8 µs downtime using coarse scheduling done by user.

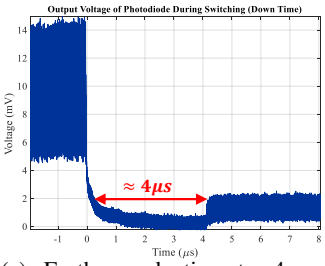

(c) Further reduction to 4 µs downtime after fine delay calibration done by user.

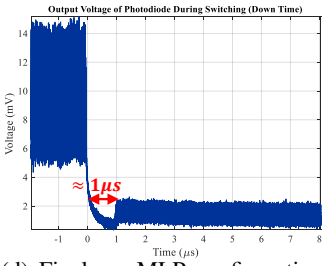

(d) Final pre-MLP configuration done by user achieving 1 µs residual downtime.

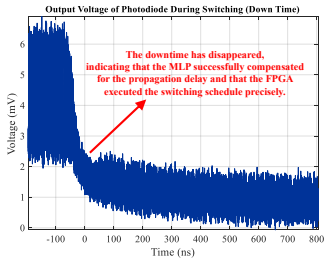

(e) Downtime eliminated after MLP compensation, confirming FPGA-scheduled switching accuracy.

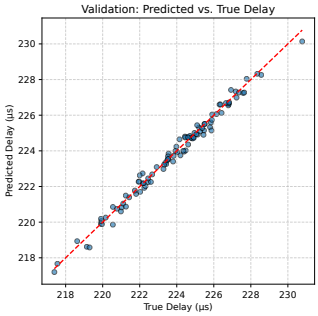

(f) Predicted vs. true delay on the validation set, confirming sub-microsecond accuracy.

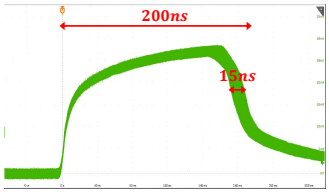

(g) System jitter measured from width variation of 200 ns short pulses.

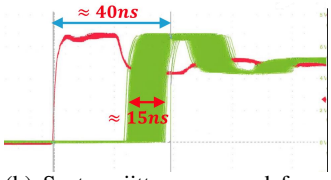

(h) System jitter measured from timing variation of the 1 PPS signal.

Fig. 3: Experimental results.

propagation delay from PTP measurements, enabling fully autonomous scheduling. The MLP model, implemented directly on FPGA, was trained using a hybrid dataset combining synthetic and real PTP delay measurements collected between two FPGA agents. A total of 10,000 samples were used, with 10% for validation and 10% for testing. As summarized in Tables I–II, the model achieved a mean absolute error below 0.3 µs and an R^2 above 0.99, accurately mapping PTP-derived delays to true fiber propagation delays over 30–70 km links. This sub-microsecond prediction accuracy directly enables precise switch-timing alignment, ensuring near-lossless optical reconfiguration during topology transitions. The validation results in Fig. 3f demonstrate the MLP's strong generalization, accurately predicting true fiber delays with sub-microsecond precision and negligible deviation from the ideal linear fit.

System performance was evaluated by transmitting 20 million packets at both 1 Gb/s and 10 Gb/s, performing optical switching midstream, and averaging results over 20 trials. As shown in Table III, the average packet loss decreased from 25.4 to 2.4 at 1 Gb/s and from 226.2 to 21.5 at 10 Gb/s—achieving nearly 90% reduction with AI-assisted scheduling. It's worth noting the existing packet loss caused by power mismatches between the two paths, which requires transponders to take time to establish new connections. At the physical layer, packet loss can be minimized to around one due to the rise/fall times of switches and clock synchronization jitter.

Photodiode voltage traces in Fig. 3 illustrate the evolution of switching precision. Without scheduling, a 250 µs outage occurs (Fig. 3a), matching the 50 km fiber delay. Manual fine-tuning reduces downtime to 8, 4, and 1 µs (Figs. 3b–3d). With AI scheduling, downtime disappears completely (Fig. 3e), confirming precise FPGA-based compensation. System jitter was further evaluated using two independent methods: (i) pulse width variation over 200 ns windows and (ii) 1 PPS rising-edge timing deviation. Both tests yielded 15 ns jitter and a 40 ns synchronization offset, shown in Figs. 3g–3h. The FPGA automatically compensates for this offset during calibration, along with path-dependent propagation mismatches between multiple fiber links. Unlike software-defined approaches, this implementation performs AI inference entirely in hardware, ensuring deterministic latency and seamless operation under network dynamics.

IV. Conclusion

We presented an FPGA-based AI-assisted optical switching system capable of real-time propagation-delay estimation and nanosecond-level scheduling. The architecture achieved 90% packet loss reduction at both 1 Gb/s and 10 Gb/s by executing an on-chip MLP model for precise delay compensation. The FPGA autonomously corrects system delays and synchronization offsets and fiber-length mismatches, ensuring deterministic, lossless reconfiguration. This work establishes a scalable foundation for self-healing, delay-aware, and AI-native optical transport networks optimized for next-generation 6G and data-center infrastructures.

Acknowledgements

This work was partially funded and supported by the Compound Semiconductor Applications Catapult and UK EPSRC project HASC (No.EP/Y037197/1).

References

[1] A. Fayad, *ICST*, 2024.
[2] C. Deepthi, *JOCN*, vol. 9, no. 2, pp. 172–182, 2017.
[3] V. Yokar, in *ONDM*. IEEE, 2025, pp. 1–3.
[4] B. Sun, in *The 51st ECOC*. IEEE, 2025.
[5] Y. Teng, in *The 51st ECOC*. IEEE, 2025.
[6] G. Patronas, *JOCN*, vol. 17, no. A87, 2025.