Event-Dropping Dynamics in Temporal Photonic Networks

Mukta Janpandit,[1,*,†] Madhuvanthi Srivatsav R,[2,†] Vadivukkarasi Jeyaselvan,[1] Shantanu Chakrabartty,[3] Chetan Singh Thakur[2] and Shankar Kumar Selvaraja[1]

[1] *Centre for Nano Science and Engineering, Indian Institute of Science, Bangalore, India*
[2] *Department of Electronic Systems and Engineering, Indian Institute of Science, Bangalore, India*
[3] *Department of Electrical and Systems Engineering, Washington University in St. Louis, St. Louis, MO, USA*
[†] *These authors contributed equally to this work.*

[*] *muktaj@iisc.ac.in*

Abstract: We experimentally demonstrate event dropping as a core computing primitive, enabling an end-to-end photonic neural network that uses delays, sorting, and dropped events for computation, achieving 97% MNIST accuracy without intermediate opto-electronic (O-E-O) conversions.

1. Introduction

Photonic delay-line neural networks have traditionally been limited to reservoir-computing paradigms, where delay-induced temporal dynamics are rich but only the readout layer is trainable. Conventional photonic neural architectures—such as MZI meshes [1] and microring-resonator crossbars [2]—depend on active optical nonlinearities or frequent O-E-O conversions, which break the photonic pipeline and incur high energy costs. Recent time-domain neuromorphic frameworks such as TEMP [3] and the processing-in-interconnect (π^2) paradigm [4] have shown that temporal primitives—programmable delay, sorting, and event dropping—can themselves serve as computational operators. Motivated by these insights, we introduce Optical Processing in Interconnect ($O-\pi^2$) paradigm, a purely time-domain photonic computing framework in which inputs are encoded as temporal events, synaptic weights are represented as optical interconnect delays, and nonlinear behavior arises intrinsically through temporal sorting and event-dropping primitives. This keeps both linear and nonlinear operations entirely within the optical domain, resulting in a streamlined end-to-end photonic inference architecture.

2. Optical Processing in Interconnect

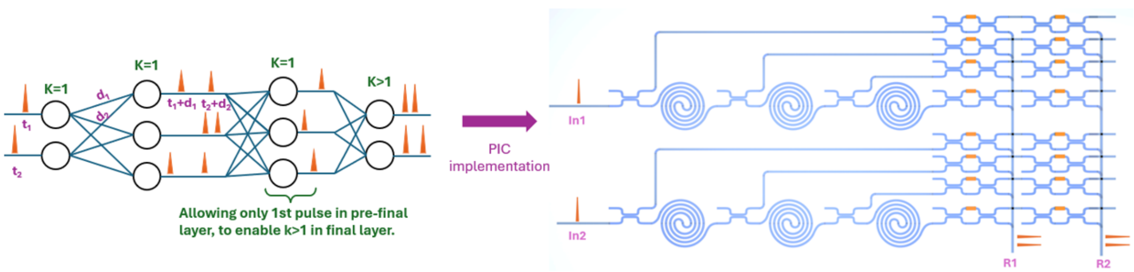

Fig. 1. Multi-layer $O-\pi^2$ architecture and its photonic implementation for a 2×2 network with 2-bit delays. An $N \times M$ network with p-bit delay quantization requires $NM2^p$ tunable couplers.

$O-\pi^2$ Networks operate entirely in the time domain, using only delay, sorting, and event dropping—formally described in [3,4]. A π^2 neuron produces an output event at time T given by: $T = \frac{M}{K} + \frac{1}{K}\sum_{j=1}^{K} t_{[j]}$, where $t_1, t_2, \ldots, t_d$ denotes the set of input event arrival times after each has been delayed by its corresponding synaptic weight, d is the fan-in, and M is a constant. The terms $t_{[1]}, \ldots, t_{[K]}$ correspond to the K earliest arrivals selected from this set. Nonlinearity arises because only the first K events contribute; later events are ignored. Setting $K = 1$ and $M = 0$ reduces the neuron to a simple first-arrival operator, which we leverage in the hidden layers to enable a fully photonic implementation, as the later events are naturally ignored by the neuron dynamics rather than explicitly dropped. For the hidden layers, $K = 1$, whereas, at the output layer, we implement $K > 1$ (after photodetection) to improve classification accuracy. Realizing this requires that pre-final layer enforce event dropping so that only the earliest K events propagate forward. Figure 1 illustrates a 2×2 $O-\pi^2$ network with 2-bit quantized weights

(4 delay lines). Inputs are time-encoded optical pulses traveling through waveguides. Delays are realized using spiral waveguides, and a tunable-coupler crossbar routes each input neuron through the appropriate delayed path to the output neuron. The outputs are time-encoded pulses that feed directly into the next layer. We experimentally

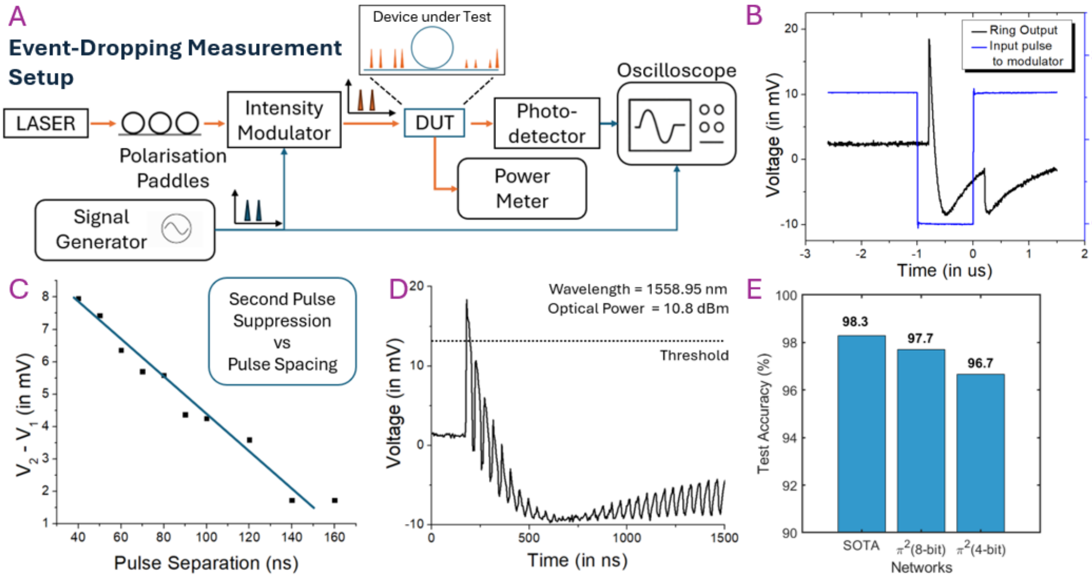

Fig. 2. Experimental demonstration of pulse suppression using a silicon microring modulator.(A) Experimental setup for characterizing nonlinear pulse suppression. (B) Measured response to a 1 µs pulse showing initial enhancement followed by rapid decay. (C) Suppression of the second pulse as a function of inter-pulse delay for 5 ns pulses. (D) Time-domain response for 5 ns pulses with 40 ns separation, illustrating event-dropping behavior. (E) MNIST test accuracy for state-of-the-art (SOTA) photonic NNs and O-PI^2 with 8-bit and 4-bit delay quantization.

demonstrate pulse suppression in a silicon photonic microring modulator to validate the event dropping mechanism for $K > 1$ layers in $O - \pi^2$ neural networks. A detailed analysis of this device's nonlinear effects was presented in earlier work [5]. Fig. 2A shows the setup used for characterization. We first launch a broad 1 µs pulse and tune both the input intensity and parking wavelength to obtain the response in Fig. 2B, where the output initially rises and then sharply decays. Parking the wavelength +0.06 nm from resonance and using 10.8 dBm input laser power yields the desired nonlinear behavior. We then inject 5 ns pulse trains with delays from 40 ns to 160 ns and measure second-pulse suppression versus delay (Fig. 2C), which determines the maximum usable delays and final-layer threshold. The first pulse induces a resonance shift via two-photon absorption and free-carrier dispersion, moving the cavity toward the pulse spectrum. Subsequent pulses arrive before the resonance recovers and thus experience reduced transmission. Figure 2D shows the measured response for a 40 ns separation. After photodetection, a simple threshold operation removes these suppressed pulses.

3. Result

Figure 2E shows the MNIST dataset accuracy of a three-layer π^2 network using trained 8-bit and 4-bit quantized delays, achieving performance close to state-of-the-art models. Crucially, we achieve this while using no electrically powered nonlinear elements in the network implementation. This distinguishes our architecture from conventional photonic neural networks that require active optical nonlinearities, electronic activation circuits, or O-E-O conversions. Our architecture naturally supports data parallelization through wavelength-division multiplexing (WDM), exploiting the inherent bandwidth of photonic links. In a multi-wavelength system, separate microring resonators perform event suppression at each wavelength; however, because we rely only on their passive nonlinear response, this scaling incurs no additional electrical power overhead. The use of microring-based nonlinearities also provides a path toward increased architectural complexity. For example, by adjusting the post-detection threshold, the network could be configured to admit the first two—or more—pulses, enabling richer temporal processing in future implementations. Optical power splitting reduces pulse energy with depth, so fully passive networks are practical only for shallow architectures unless optical gain is introduced.

References

1. Bandyopadhyay, S., Sludds, A., Krastanov, S. et al. Single-chip photonic deep neural network with forward-only training. Nat. Photon. 18, 1335–1343 (2024).
2. Feldmann, J., Youngblood, N., Karpov, M. et al. Parallel convolutional processing using an integrated photonic tensor core. Nature 589, 52–58 (2021).
3. R. M. Srivatsav, S. Chakrabartty and C. S. Thakur, "Neuromorphic Computing With Address-Event-Representation Using Time-to-Event Margin Propagation," in IEEE Journal on Emerging and Selected Topics in Circuits and Systems, vol. 13, no. 4, pp. 1114-1124, Dec. 2023.
4. M. Srivatsav., Bhattacharyya, C., Chakrabartty, S., Thakur, C. S. (2025). When Routers, Switches and Interconnects Compute: A processing-in-interconnect Paradigm for Scalable Neuromorphic AI. arXiv preprint arXiv:2508.19548.
5. V. Jeyaselvan and Selvaraja, S. K., "Mitigation of carrier induced optical bistability in silicon ring resonators", Optics Communications, vol. 493, p. 127021, 2021.
6. L. Deng, "The MNIST Database of Handwritten Digit Images for Machine Learning Research [Best of the Web]," in IEEE Signal Processing Magazine, vol. 29, no. 6, pp. 141-142, Nov. 2012.