Research Square

Preprints are preliminary reports that have not undergone peer review.
They should not be considered conclusive, used to inform clinical practice,
or referenced by the media as validated information.

Scalable Microring-Based Interconnects with Channel Equalization and Closed-Loop Stabilization

Yating Wan

yating.wan@kaust.edu.sa

Integrated Photonics Lab, King Abdullah University of Science and Technology (KAUST)
https://orcid.org/0000-0003-2157-2406

Zichun Lin
Integrated Photonics Lab, King Abdullah University of Science and Technology (KAUST)

Zhican Zhou
King Abdullah University of Science and Technology

Xiangpeng Ou
King Abdullah University of Science and Technology

Weiping Li
Integrated Photonics Lab, King Abdullah University of Science and Technology (KAUST)

Xuhao Wu
Integrated Photonics Lab, King Abdullah University of Science and Technology (KAUST)

Hongyan Yu
King Abdullah University of Science and Technology

Xiaoyi Song
Integrated Photonics Lab, King Abdullah University of Science and Technology (KAUST)

William He
Integrated Photonics Lab, King Abdullah University of Science and Technology (KAUST)

Article

Keywords:

Posted Date: July 10th, 2026

DOI: https://doi.org/10.21203/rs.3.rs-10269988/v1

License: This work is licensed under a Creative Commons Attribution 4.0 International License.
Read Full License

Additional Declarations: There is no conflict of interest

Scalable Microring-Based Interconnects with Channel Equalization and Closed-Loop Stabilization

Zichun Lin[1†], Zhican Zhou[1†], Xiangpeng Ou[1], Weiping Li[1], Xuhao Wu[1], Hongyan Yu[1], Xiaoyi Song[1], William He[1], Yating Wan[1*]

[1*]Integrated Photonics Lab, CEMSE, King Abdullah University of Science and Technology, Thuwal, 23955, Makkah Region, Saudi Arabia.

*Corresponding author(s). E-mail(s): yating.wan@kaust.edu.sa;
†These authors contributed equally to this work.

Abstract

Optical interconnects are essential for scaling photonic computing from isolated cores to distributed multi-core systems. In analog photonic processors, however, the interconnect must preserve multiwavelength optical power during inter-core signal transfer, rather than simply reconfigure optical paths. Conventional switching fabrics do not compensate channel-power imbalance, wavelength-dependent loss, or resonance drift, allowing these impairments to accumulate as errors across cascaded computations. Here, we propose and experimentally demonstrate a photonic interconnect unit (PIU) that integrates a 4-by-4 microring-resonator switching fabric, cascaded programmable channel equalizers, and a three-phase field-programmable gate array-assisted feedback controller on a 180-nm silicon photonic platform. The PIU supports reconfigurable wavelength-division-multiplexed redistribution, including multiplexing and demultiplexing, while cascaded equalization, stabilized by closed-loop feedback, suppresses residual channel-power variation to below 0.2 dB. The feedback controller achieves 10.3-bit-equivalent steady-state control precision and maintains signal-to-noise ratio fluctuation below 0.5 dB across a 22 °C temperature variation. Eight-channel multiwavelength optical weight reconstruction reduces the normalized end-to-end error from 0.179 to 0.0193 (9.3-fold). For ResNet-18, hardware-aware inference simulations under a measurement-derived 6.5-dB channel-mismatch model show that equalization recovers network accuracy from 54.33% to the clean hardware-reference level of approximately 94.6%. These results establish

actively stabilized photonic interconnects as an analog signal-conditioning layer for scalable multi-core photonic computing clusters.

1 Introduction

Photonic computing has emerged as a high-throughput and energy-efficient platform for artificial intelligence (AI) acceleration [1–5]. Realizing the full potential of this platform, however, requires moving beyond isolated single-core processors toward multi-core architectures in which multiple photonic computing units (PCUs) are interconnected to form scalable photonic computing clusters[6, 7]. The critical missing element is an interconnect layer capable of transferring intermediate computational states between PCUs entirely in the optical domain while preserving analog amplitude fidelity. This requirement is fundamentally more demanding than those imposed by digital optical links[8, 9].

Digital optical links can tolerate moderate wavelength-dependent loss and path-dependent insertion loss provided that logic-state margins remain sufficient. In contrast, analog photonic processors use optical power or field amplitude as computational variables [10]. As depicted in Fig. 1(a), wavelength-dependent loss, path-dependent insertion loss, and resonance drift can produce inter-channel power variations, which distort analog optical signals and propagate as computational errors across cascaded layers[11], especially in deep neural-network workloads, as schematically illustrated in Fig. 1(b). This challenge is further exacerbated in microring-resonator (MRR)-based interconnects, whose resonance conditions are highly sensitive to fabrication variations, thermal crosstalk, and environmental fluctuations[12, 13].

Prior MRR-control studies have made important progress toward reliable photonic computing, including calibrated MRR weight-bank operation in microcomb-driven photonic processors [14], high-precision microring synaptic weight programming [10], self-calibrating microring synapses under crosstalk and environmental perturbations [15, 16], and individual-resonance or weight-function stabilization using wavelength-locking and feedback-control schemes [17, 18]. These studies have established important device- and core-level foundations, but they do not directly address the interconnect-level signal-preservation problem that appears when multiwavelength analog optical states are routed between distributed PCUs. In this regime, wavelength-selective routing is both a connectivity function and a source of the wavelength- and path-dependent power imbalance, while resonance drift can perturb both the selected route and the equalization state. Preserving analog fidelity at the interconnect level therefore requires wavelength-selective routing, post-routing channel-power equalization, and closed-loop resonance stabilization to be co-designed within the same photonic module.

Here, we address these challenges by proposing and experimentally demonstrating a photonic interconnect unit (PIU) for scalable distributed photonic computing. The PIU integrates a compact 4×4 MRR-based wavelength-selective switching fabric, programmable equalization MRRs, and field-programmable gate array (FPGA)-assisted

feedback control. The switching fabric uses a compact triangular MRR mesh to support reconfigurable wavelength division multiplexing (WDM) signal redistribution, including wavelength-selective multiplexing and demultiplexing.

After routing, cascaded equalization stages act as programmable wavelength-selective attenuators that trim stronger channels toward a common target power without optical–electrical–optical conversion. Experimentally, the fabricated silicon photonic PIU demonstrates reconfigurable WDM signal redistribution with crosstalk suppression exceeding 20 dB for selected DEMUX channels, and two-stage equalization suppresses residual multiwavelength channel-power variation to below 0.2 dB.

To stabilize the MRR operating points, we implement a three-phase FPGA-assisted proportional–integral–derivative (PID) locking strategy consisting of voltage–power lookup-table construction, monotonic target acquisition, and real-time feedback locking. Under the implemented control conditions, the controller converges within 40 ms, provides a sourcemeter-characterized 10.3-bit-equivalent steady-state control precision, and maintains stable transmission with signal-to-noise ratio fluctuation within 0.5 dB over a 22 °C temperature variation.

Finally, we evaluate the computing impact of the proposed PIU through multi-wavelength optical weight reconstruction and hardware-aware inference simulations on VGG-11 and ResNet-18. Eight-channel optical weight reconstruction achieves a 9.3-fold reduction in normalized end-to-end reconstruction error, with root-mean-square error (RMSE) reduced from 0.179 to 0.0193 under closed-loop stabilization. In hardware-aware ResNet-18 CIFAR-10 simulations under a measurement-derived 6.5-dB channel-mismatch distribution, inference accuracy was recovered from 54.33% without equalization to the clean hardware-reference level of approximately 94.6% after equalization.

Together, these results establish the PIU as an actively stabilized analog signal-conditioning layer, rather than a passive optical routing fabric. By co-designing routing reconfigurability, post-routing equalization, and drift-tolerant stabilization in a single unit, the PIU directly addresses a key scaling bottleneck in distributed photonic computing: intermediate optical signals must remain in the analog domain while maintaining sufficient amplitude fidelity for downstream computation.

2 Results

2.1 PIU architecture and operating methods

The PIU introduced in Fig. 1(c) is implemented as a two-stage wavelength-selective photonic circuit consisting of an MRR-based switching fabric followed by programmable equalization MRRs. Multiwavelength analog optical signals transmitted from an upstream PCU first enter the switching fabric for reconfigurable wavelength routing and are then processed by the equalization stage for channel-power balancing before being delivered to the downstream PCU. The device was fabricated on a 180-nm silicon photonic platform using 10 μm-radius MRRs, providing the wavelength selectivity and thermo-optic tunability required for WDM routing and closed-loop channel-power control.

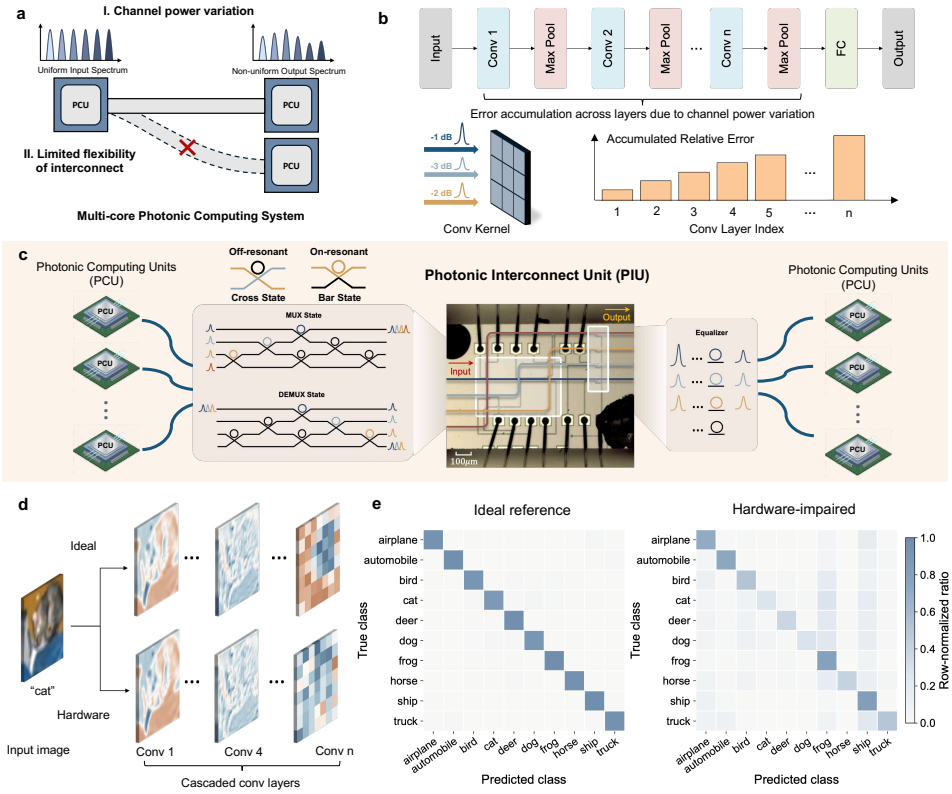


Fig. 1 Photonic interconnect challenges in multi-core photonic computing and the proposed photonic interconnect unit (PIU). (a) Two key challenges in multi-core photonic computing systems: channel-power variation during inter-PCU transmission and limited routing flexibility between photonic computing units. (b) Schematic illustration of error accumulation in cascaded convolutional layers caused by channel-dependent power variation. (c) Architecture of the proposed PIU, which integrates an MRR-based wavelength-selective MUX/DEMUX switching fabric, a fabricated photonic chip, and programmable equalization MRRs for reconfigurable inter-PCU optical signal transfer. (d) Representative feature maps under clean-reference operation and an illustrative hardware-impaired interconnect condition with channel-power variation. (e) Row-normalized confusion matrices of inference under the corresponding clean-reference and illustrative hardware-impaired conditions.

Programmable photonic routing and linear transformation have been widely implemented using mesh-based architectures. Universal interferometric meshes, such as the Reck and Clements architectures, provide systematic decompositions for arbitrary linear or unitary transformations [19, 20]. In parallel, optical switching fabrics based on crossbar, Beneš [21], Banyan/Omega-type [22], and related multistage topologies have been extensively studied for reconfigurable optical connectivity, with different tradeoffs among blocking behavior, switch count, path-dependent loss, crosstalk, and control complexity [23, 24]. In this work, the routing stage is not intended to serve as a universal unitary processor or a strictly nonblocking digital switch. Instead, it is

designed as a compact wavelength-selective fabric for reconfigurable WDM MUX/DEMUX and inter-PCU analog signal transfer. Accordingly, we adopt a triangular mesh of cascaded 2×2 MRR crossbar elements, which provides the required wavelength-selective routing functionality with reduced hardware and control complexity. In the off-resonant state, each MRR crossbar element operates in the cross state, whereas tuning the MRR resonance to a target wavelength switches that wavelength channel to the bar state. Six MRR crossbar elements are cascaded to form a compact 4×4 wavelength-selective switching fabric. Compared with a dense 4×4 crossbar requiring 16 wavelength-selective crosspoints, the triangular fabric uses only six MRR crossbar elements, reducing the number of routing elements and control variables by 62.5%. Despite this reduced hardware complexity, the fabric supports flexible wavelength-selective routing functions. The generalized routing capability for an $N \times N$ fabric with K wavelength channels is summarized in Supplementary Table 1.

Although the switching fabric provides the required routing reconfigurability, wavelength routing alone cannot guarantee accurate analog signal transfer between PCUs. Routing through different wavelength-selective paths can introduce channel-dependent transmission imbalance due to wavelength-dependent loss, resonance detuning, and path-dependent insertion loss. For amplitude-encoded photonic computing, such channel-power nonuniformity directly scales the optical signals delivered to the next computing unit and can therefore perturb the effective input channels of subsequent convolutional operations.

Fig. 1(d,e) conceptually illustrates the computational impact of channel-power mismatch in downstream neural-network inference. Compared with the clean reference, the hardware-impaired case exhibits visibly distorted feature maps across cascaded convolutional layers and degraded classification behavior in the row-normalized confusion matrices. These results motivate the need for an interconnect that not only provides wavelength-selective routing, but also compensates channel-power nonuniformity to preserve analog signal fidelity between PCUs.

To provide this compensation, the routed optical signals are subsequently sent to the equalization MRRs, as illustrated in Fig. 1(c). Each equalization MRR functions as a programmable wavelength-selective attenuator. Previous MRR-based photonic weighting schemes have used coupled- or higher-order MRR structures to shape the spectral response of weight elements [25]. For example, dual-coupled MRRs in an add-drop configuration have been employed to realize flatter passbands and broader optical bandwidth for high-speed optical weighting [26]. In this work, the equalization stage serves a different function. The equalization MRR, cascaded after the routing crossbar, introduces a tunable resonance notch onto the wavelength-selected transmission peak formed by the routing stage. By thermally detuning this notch with respect to the routed channel peak, the spectral overlap between them can be continuously adjusted, thereby controlling the amount of notch-induced attenuation applied to that wavelength channel. This mechanism allows higher-power wavelength channels to be selectively trimmed toward a common target level while keeping the intermediate analog signals in the optical domain. As a result, the cascaded routing-and-equalization architecture enables the PIU to aggregate, redistribute, and deliver multiwavelength analog optical signals while compensating interconnect-induced channel-power imbalance before transmission to downstream PCUs.

2.2 Experimental characterization of wavelength routing and channel equalization

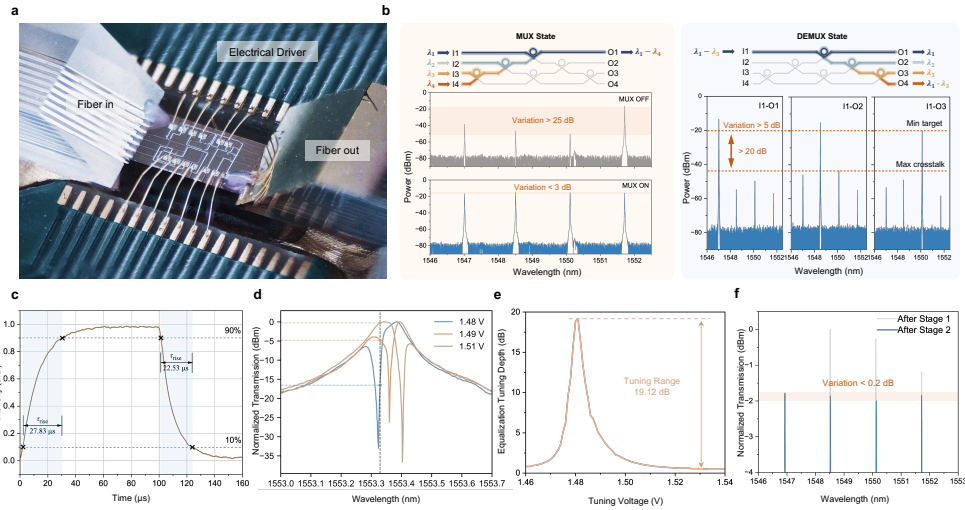


Fig. 2 Experimental characterization of the fabricated PIU, including wavelength-selective routing, switching dynamics, and channel equalization. (a) Photograph of the packaged PIU chip with fiber input/output coupling and electrical heater-driving probes. (b) Measured wavelength-routing characteristics of the MRR switching fabric in the MUX state and DEMUX state. (c) Temporal switching response of the routing unit, showing microsecond-scale rising and falling transitions under electrical heater actuation. (d) Measured transmission spectra of an equalization MRR under different heater voltages. (e) Extracted equalization tuning depth as a function of heater voltage, showing a tuning range of 19.12 dB. (f) Measured multiwavelength channel powers after routing stage and equalization stage.

We first experimentally characterize the wavelength-routing, switching, and channel-equalization functions of the fabricated PIU. Fig. 2(a) shows the packaged chip used in the measurement, with optical fiber input/output coupling and wire-bonded electrical connections for heater control.

The wavelength-routing behavior is characterized by measuring the transmission spectra under representative MUX and DEMUX configurations, as shown in Fig. 2(b). In the MUX configuration, four wavelength channels from different input ports are routed and combined toward a common output port. In the unconfigured state, the selected output exhibits a channel-power spread exceeding 25 dB because the target wavelengths are not simultaneously routed to the common output. After configuring the MUX state, the four selected wavelength channels are routed to output port 1 with a residual power spread below 3 dB. The measured fiber-to-fiber insertion losses for the four routed MUX channels are 16.46, 14.67, 14.94, and 16.16 dB for λ_1–λ_4, respectively.

In the DEMUX configuration, multiwavelength signals launched from input port 1 are wavelength-selectively separated and directed to different output ports. The measured spectra confirm port-dependent wavelength routing, with fiber-to-fiber insertion losses of 13.46, 15.17, and 19.96 dB for the selected $I1$-to-$O1$, $I1$-to-$O2$, and $I1$-to-$O3$ channels, respectively. The non-target wavelength components remain suppressed, giving crosstalk suppression exceeding 20 dB for the selected DEMUX channels.

The temporal switching response of the routing unit is shown in Fig. 2(c). The measured optical intensity follows the electrical heater actuation with a 10–90% rise time of 27.83 μs and a 90–10% fall time of 22.53 μs. This microsecond-scale switching response confirms that the MRR routing state can be electrically reconfigured on a timescale compatible with dynamic interconnect operation.

To compensate for routing-induced power imbalance, the routed signals are subsequently processed by the equalization MRRs. Each equalization MRR operates as a programmable wavelength-selective attenuator that trims excess power from stronger wavelength channels. Fig. 2(d) experimentally verifies this attenuation mechanism by showing representative transmission spectra of an equalization MRR under different heater voltages. Changing the heater voltage shifts the through-port resonance notch relative to the selected operating wavelength, thereby producing a pronounced change in the transmitted power at that wavelength.

The extracted equalization tuning depth is shown in Fig. 2(e). The equalization MRR provides a tuning depth of 19.12 dB over the measured voltage range, offering sufficient dynamic range to attenuate stronger wavelength channels and compensate routing-induced channel-power imbalance. Fig. 2(f) further shows the measured multiwavelength channel powers before and after equalization. After the routing stage, the channel powers remain non-uniform because of wavelength- and path-dependent transmission differences. After the equalization stage, the stronger channels are trimmed toward the common target level, and the residual channel-power variation is suppressed to below 0.2 dB, as indicated by the shaded band.

These measurements collectively demonstrate the three core hardware capabilities of the fabricated PIU, reconfigurable wavelength-selective routing, microsecond-scale electrical switching, and high-dynamic-range channel equalization. Together, these capabilities establish the hardware basis for stable analog optical signal transfer between cascaded PCUs while suppressing routing-induced channel-power nonuniformity.

2.3 FPGA-based feedback stabilization for accurate and robust operation

Although the routing and equalization results demonstrate the static functionality of the PIU, accurate analog optical signal transfer requires the MRR operating points to be actively maintained during operation. MRR resonances are sensitive to temperature fluctuations, fabrication-induced mismatch and thermal crosstalk, so open-loop voltage biasing can lead to power drift and degrade the equalized channel states. We therefore incorporate a three-phase FPGA-assisted feedback locking (TP-FAPL) scheme into the PIU, forming a photonic–electronic co-designed stabilization loop for real-time resonance tracking and channel-power locking.

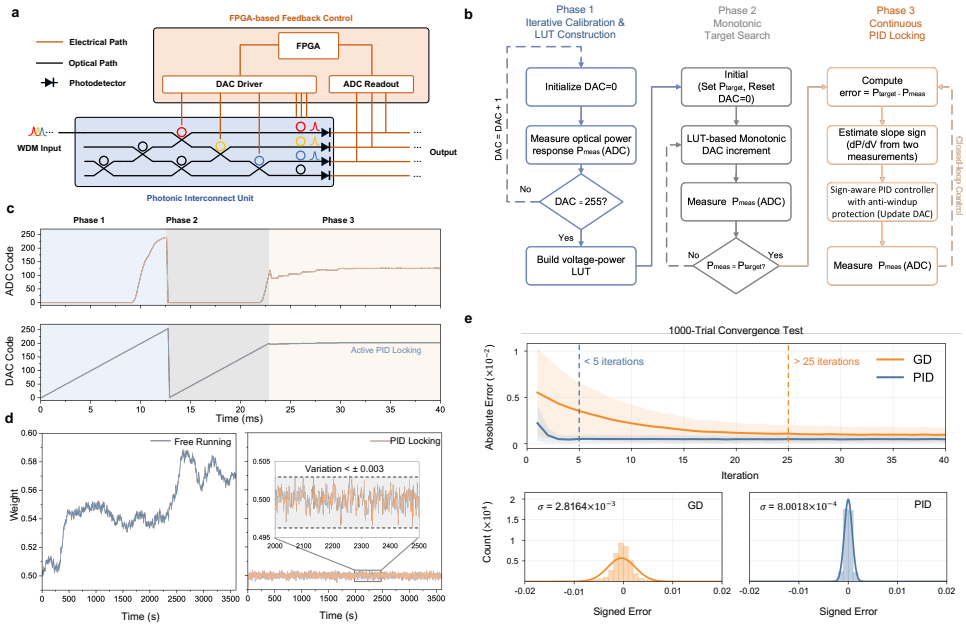


Fig. 3 FPGA-based feedback control for resonance stabilization and channel equalization. (a) Experimental schematic of the feedback-controlled PIU system. (b) TP-FAPL control algorithm consisting of voltage–power lookup table (LUT) construction, monotonic target search, and PID locking. (c) Temporal evolution of the ADC and DAC signals during the control sequence. (d) Long-term stability over 1 h with free running and PID locking. (e) 1000-trial convergence test of gradient-descent (GD) optimization and PID control.

Fig. 3(a) shows the schematic of the feedback-controlled PIU system. The optical power of the selected wavelength channel is monitored and converted into an electrical feedback signal, which is processed by a custom FPGA-based controller. The controller compares the measured power with the target value and continuously updates the heater-driving voltage, thereby compensating resonance drift and stabilizing the transmission of the target channel. This closed-loop control path enables the programmed optical state to be maintained at the hardware level without offline post-processing.

As depicted in Fig. 3(b), TP-FAPL follows a three-phase sequence: voltage–power lookup-table (LUT) construction, monotonic target acquisition, and sign-aware PID locking. During PID locking, the controller estimates the local slope sign of the MRR transfer curve and uses a sign-aware power error to preserve the correct feedback direction on different resonance branches. The digital-to-analog converter (DAC) code is updated with bounded integral accumulation and DAC-range saturation to avoid excessive control excursions. The temporal analog-to-digital converter (ADC) and DAC traces in Fig. 3(c) show the transition from target acquisition to active PID locking. The corresponding real-time eye-diagram evolution during the three-phase locking sequence is provided in Supplementary Video 1. Further implementation details are provided in Materials and methods and Supplementary Algorithm 1.

We next evaluate the long-term stability of the PIU under free-running and PID-locked conditions. As shown in Fig. 3(d), the free-running output weight exhibits substantial drift over the 1-hour measurement, increasing from 0.50 to nearly 0.59, with a total drift approaching 0.1 in normalized weight. This drift is consistent with ambient thermal fluctuations and the resulting resonance shift of the MRR under fixed heater bias. In contrast, when PID locking is enabled, the output weight remains tightly stabilized around the target value over the same measurement period. The residual normalized-weight fluctuation is suppressed to within ± 0.003, as shown in the inset. These results confirm that FPGA-based feedback control is effective for maintaining stable channel equalization during long-term operation.

Finally, we compare the control behavior of gradient-descent (GD) and PID locking using a 1000-trial convergence test with randomized initial optical-weight states, as shown in Fig. 3(e). The upper panel plots the mean absolute error as a function of iteration, with the shaded regions indicating the mean $\pm$ standard deviation over the 1000 trials. Under the same implemented control protocol, PID locking rapidly reduces the error and reaches the steady-state error band within fewer than 5 iterations, whereas GD requires more than 25 iterations to approach a comparable steady-state region.

The steady-state control precision is quantified using the bit-precision metric

$$\text{Bit precision} = \log_2 \left(\frac{w_{\max} - w_{\min}}{\text{RMSE}_w} \right), \tag{1}$$

where $w_{\max} - w_{\min} = 1$ for normalized optical weights[10]. Because the residual signed errors are centered near zero, the standard deviation of the fitted error distribution approximates RMSE_w. To avoid the resolution limit of the 8-bit ADC in the FPGA readout path, the steady-state error distributions were further characterized using a high-resolution sourcemeter. The measured standard deviations are 2.8164×10^{-3} for GD and 8.0018×10^{-4} for PID. According to Eq. (1), these values correspond to bit-equivalent precisions of 8.5 and 10.3 bits, respectively. GD-based control reduces the power error through iterative voltage updates, but its convergence is sensitive to the chosen step size, measurement noise, and the local slope of the nonlinear MRR transfer curve. After a suitable monotonic operating region is identified, PID locking directly regulates the measured power error through closed-loop feedback, enabling continuous compensation of resonance drift and suppression of residual power fluctuations.

2.4 Temperature-robust transmission under environmental perturbations

Temperature variation is a major source of instability in MRR-based photonic interconnects because the thermo-optic effect shifts the resonance wavelength and changes the transmission of the selected wavelength channel. Such drift directly perturbs the analog optical power delivered to the downstream PCU and can degrade the fidelity of inter-PCU signal transmission. To evaluate the temperature robustness of the proposed PIU, we measure eye diagrams under controlled temperature variations using the experimental setup shown in Fig. 4(a). A thermoelectric cooler (TEC) controller is used to apply controlled chip-level temperature perturbations during operation by

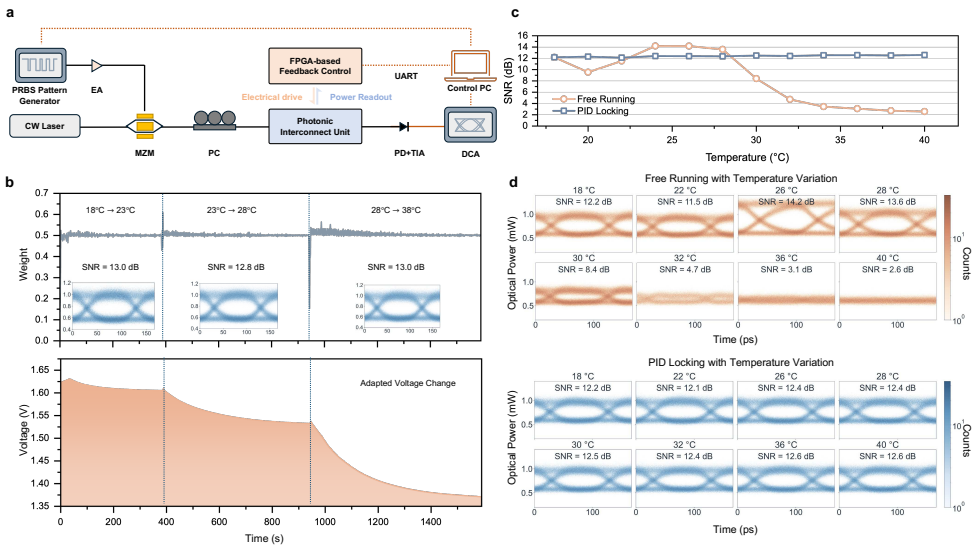


Fig. 4 Transmission stability under temperature variation with closed-loop PID stabilization. (a) Experimental setup for eye-diagram measurement. (b) System response under staged temperature variation, including optical weight, eye diagrams, and heater-voltage adaptation. (c) Measured signal-to-noise ratio (SNR) as a function of temperature for free-running and PID-locked operation. (d) Representative eye diagrams at different temperatures with and without PID locking.

adjusting the substrate temperature of the photonic chip. A 10-Gbaud non-return-to-zero (NRZ) pseudo-random binary sequence (PRBS) generated by a pattern generator is amplified by an electrical amplifier and used to drive the Mach–Zehnder modulator. The resulting intensity-modulated optical signal is transmitted through the PIU, detected by a photodetector and transimpedance amplifier, and recorded by a digital communication analyzer (DCA) for eye-diagram characterization.

Fig. 4(b) shows the closed-loop response under staged temperature variations. Here, the optical weight is defined as the normalized transmitted optical power mapped to the target analog weight value. As the temperature is varied over the staged intervals of 18–23 °C, 23–28 °C, and 28–38 °C, the PID controller continuously adjusts the heater voltage to compensate for the thermally induced resonance shift. The required heater voltage decreases with increasing ambient temperature, consistent with the red-shift of the MRR resonance under thermal perturbation. As a result, the optical weight remains close to the target value, and the representative eye diagrams maintain clear openings with SNR values of 13.0, 12.8, and 13.0 dB for the three temperature intervals.

Fig. 4(c) and 4(d) compare free-running and PID-locked transmission over a wider temperature range from 18 to 40 °C. In the free-running case, the heater voltage is fixed after initial alignment, so resonance drift causes strong SNR variation and progressive eye closure as the chip temperature changes. The SNR remains above 10 dB at lower temperatures but decreases rapidly at elevated temperatures, falling to 8.4 dB at 30 °C, 4.7 dB at 32 °C, and 2.6 dB at 40 °C. In contrast, PID locking maintains

the SNR within 12.1–12.6 dB across the full temperature range, corresponding to an SNR fluctuation of only 0.5 dB over a 22 °C temperature variation. The representative eye diagrams confirm that closed-loop stabilization preserves the transmission quality of the 10-Gbaud NRZ signal under thermal perturbations. These results show that closed-loop stabilization can maintain the target operating point during temperature changes, preserving the 10-Gbaud transmission quality without re-alignment.

2.5 Multiwavelength optical weight reconstruction



Fig. 5 Multiwavelength optical weight reconstruction using closed-loop stabilization. A grayscale KAUST pattern is encoded as a normalized optical-weight matrix across eight wavelength channels. The rows represent wavelength channels, and the columns represent sequential weight samples. After LUT calibration, the target weights are reconstructed under free-running and PID-stabilized operation. The right panels compare the measured weights with the target values. PID stabilization substantially suppresses the end-to-end reconstruction error, reducing the normalized root-mean-square error (RMSE) from 0.179 to 0.0193.

While the temperature-dependent eye-diagram measurements verify stable signal transmission under environmental perturbations, analog photonic computing further requires accurate preservation of continuous optical power levels. In such systems, optical powers represent analog weights or intermediate computational values; therefore, channel-dependent transmission errors and resonance drift can cause the measured optical weights to deviate from their target values and lead to accumulated errors in cascaded photonic systems. To directly evaluate the analog optical-weight fidelity of the proposed PIU, we encode a grayscale KAUST pattern into a target optical-weight matrix and reconstruct it across eight wavelength channels.

The target pattern is first converted into a normalized weight matrix with values between 0 and 1. The rows of the matrix correspond to the eight wavelength channels used for WDM optical-weight mapping, while the columns represent sequential weight samples. After LUT calibration, each target weight is mapped to a corresponding heater voltage and optical transmission state. In the free-running case, the calibrated voltages are applied without feedback correction. In the PID-stabilized case, the transmitted optical power is continuously monitored and locked to the corresponding target

Table 1 Comparison of TP-FAPL with representative MRR-control approaches.

Method	Reported precision	Locking time	Temp. robustness
Self-homodyne locking [17]	–	75 ms + ∼10 s	–
Dithering control [10]	9.0-bit	13–16 s	–
Dual-wavelength synchronization [15]	7.0-bit	1 s	0.5 °C
Gradient-descent control [14]	9.0-bit	–	–
Single-monitor calibration [16]	4.5-bit	–	0.2 °C
Self-referenced PID [18]	9.3-bit	–	6 °C
Modulation-and-weighting MRR neuron [27]	9.0-bit	–	0.5 °C
TP-FAPL, this work	**10.3-bit**[a]	**< 40 ms**	**>20 °C**

[a]The 10.3-bit precision of this work was characterized using a high-resolution sourcemeter to avoid the resolution limit of the FPGA ADC readout.

value through PID control. The reconstructed weight matrices and target-output scatter plots are then used to quantify the multiwavelength analog weight-mapping error.

As shown in Fig. 5, the free-running reconstruction exhibits evident intensity distortion and channel-dependent nonuniformity, indicating that LUT calibration alone is insufficient to maintain accurate optical weights under drift. In contrast, PID-stabilized operation reconstructs the target pattern with substantially improved fidelity, and the measured weights exhibit reduced deviation from the ideal $y = x$ reference line. Quantitatively, the normalized end-to-end reconstruction RMSE is reduced from 0.179 to 0.0193, corresponding to a 9.3-fold error reduction. This RMSE evaluates the complete multiwavelength reconstruction process across all target weight levels and wavelength channels, and is therefore distinct from the single-state steady-state control precision characterized in Fig. 3(e).

Table 1 compares the proposed TP-FAPL strategy with representative MRR-control approaches in terms of reported precision, locking time, and temperature robustness. Since these works address different MRR-based functions, the comparison is intended to highlight representative control capabilities rather than provide a one-to-one benchmark. TP-FAPL provides a combination of 10.3-bit-equivalent control precision, sub-40-ms FPGA-assisted locking, and stable operation under temperature variations exceeding 20 °C. For photonic interconnects, this capability is important because the locked state must remain stable after routing and equalization, rather than only during the initial optical-state setting.

2.6 System-level scalability and inference robustness

The experimental results above validate the ability of the proposed PIU to provide reconfigurable wavelength routing, channel-power equalization and closed-loop stabilization at the interconnect level. To assess how these device-level functions

can support scalable photonic computing, Fig. 6(a) illustrates a conceptual distributed photonic computing architecture based on the proposed PIU. Multiple PCUs are organized into photonic compute clusters, and different clusters are connected through a reconfigurable spine-level optical trunk interconnect [28]. In this architecture, source-side PIUs aggregate and route WDM signals from upstream PCUs, while destination-side PIUs demultiplex and redistribute selected wavelength channels toward downstream PCUs. Shared semiconductor optical amplifier (SOA) blocks can be inserted after the source-side PIUs to provide coarse global power recovery of the WDM bundle [29], whereas the PIU equalization stage compensates wavelength- and route-dependent residual channel-power imbalance. The high-level control plane performs workload mapping, wavelength assignment and routing configuration according to the computing task, while local PIU controllers configure MUX/DEMUX routing states, maintain the optical operating points and stabilize the target channel powers.

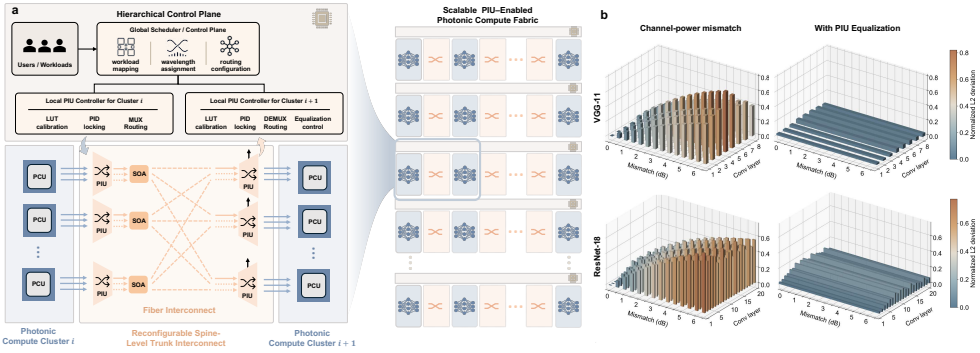


Fig. 6 Scalable photonic interconnect architecture enabled by the proposed PIU. (a) Conceptual architecture of distributed photonic compute clusters connected through a reconfigurable spine-level optical trunk interconnect. (b) Simulated layer-wise relative L_2-norm output error in VGG-11 and ResNet-18 under different unequalized channel-power mismatch levels and after PIU equalization.

Based on this scalable PIU architecture, we further evaluate how interconnect-induced channel-power mismatch affects network-level computational fidelity using hardware-aware simulations with two representative convolutional neural-network architectures (VGG-11 and ResNet-18) on CIFAR-10, a standard ten-class image-classification dataset [30–32]. The networks were trained under the clean reference condition and then evaluated under different channel-power mismatch levels. For the unequalized mismatch cases, the channel-loss classes were derived from the measured routed-channel powers in Fig. 2(b). For the equalized case, the residual channel-power losses were bounded between 0 and 0.2 dB, consistent with the measured post-equalization variation in Fig. 2(f). Detailed simulation settings are provided in Materials and methods and Supplementary Information.

As shown in Fig. 6(b), the measurement-derived channel-power mismatch model induces substantial feature-map deviations across the convolutional layers. For VGG-11, under the strongest 6.5-dB unequalized mismatch condition, the maximum

layer-wise relative L_2-norm error reaches 0.8208 without equalization and is reduced to 0.0597 after equalization. This suppression of feature-map error also preserves inference accuracy: the CIFAR-10 accuracy decreases from the clean hardware-reference value of 92.92% to 66.97% without equalization, whereas equalization restores it to approximately 92.9%. For ResNet-18, the mismatch-induced distortion is distributed across more convolutional layers and residual branches. Under the same 6.5-dB mismatch condition, the maximum layer-wise relative L_2-norm error reaches 0.7728 without equalization and is suppressed to 0.1413 with equalization. Correspondingly, the ResNet-18 accuracy drops from the clean hardware-reference value of 94.63% to 54.33% without equalization, while equalization restores it to approximately 94.6%. These results indicate that PIU equalization can substantially suppress the network-level degradation induced by interconnect channel-power mismatch and recover near-clean inference accuracy under the simulated hardware conditions.

3 Discussion

The experimental demonstrations and hardware-aware simulations establish the PIU as an active analog signal-conditioning layer, rather than merely a passive routing fabric. This distinction is important in power-encoded photonic computing, where the interconnect is not only a connectivity resource but also a potential source of computational error. The wavelength-selective routing process that redistributes signals between cores can introduce wavelength- and path-dependent power imbalance, which then propagates through cascaded downstream computations. By co-locating routing, post-routing equalization, and closed-loop stabilization in a single module, the PIU compensates this imbalance near the point where it is introduced, rather than relying on electronic post-processing or correction within the computing cores.

Specifically, the PIU fulfills three interconnect-level functions required for distributed photonic computing: wavelength-selective WDM signal redistribution, channel-power equalization, and closed-loop resonance stabilization. The MRR-based switching fabric enables reconfigurable routing, MUX, and DEMUX operations; the equalization stage suppresses wavelength- and path-dependent power imbalance after routing; and the FPGA-assisted feedback loop stabilizes the programmed optical states against resonance drift and temperature variation.

At the system level, the proposed PIU can be extended from a single interconnect node to a hierarchical photonic compute cluster, in which a global scheduler can assign workloads [33], wavelengths and routing states [34], while distributed local PIU controllers maintain the optical operating points of individual routing and equalization nodes.

The scalability of the control plane is a first practical consideration for large MRR-based photonic systems. A straightforward one-controller-per-MRR implementation can become area-inefficient because the electronic wavelength-locking controller may occupy a footprint much larger than the MRR itself, an overhead that increases as more sophisticated control algorithms are adopted [35]. This does not, however, preclude FPGA-assisted feedback control because digital control logic operates on a much

faster timescale than the thermo-optic response of MRRs, which is naturally suitable to shared, pipelined or time-division-multiplexed stabilization. Exploiting this speed mismatch between the heater dynamics and the electronic controller, pipelined time-division-multiplexing control has been shown to maintain wavelength locking for multiple MRRs with a shared controller while preserving the effective loop bandwidth [35]. The FPGA-assisted controller demonstrated here should therefore be viewed as a flexible prototype platform for validating the feedback strategy, while larger PIU networks map the same control principle onto hierarchical, shared or integrated electronic–photonic implementations.

A second consideration is that shared optical gain introduces complementary trade-offs. An SOA or erbium-doped fiber amplifier (EDFA) can provide coarse global power recovery for a WDM bundle and reduce the impact of common insertion loss, but does not, by itself, correct wavelength-dependent channel imbalance [36, 37]. Moreover, non-ideal gain flatness, amplified spontaneous emission, gain saturation and inter-channel gain dynamics can add noise and distortion in large-scale systems [38]. Shared optical gain should therefore be regarded as coarse power-budget recovery, with the PIU equalization stage remaining necessary for fine, wavelength-selective channel-power balancing. Future implementations may combine on-chip monitor photodiodes, integrated driver electronics, shared digital controllers and optical gain management to reduce control overhead and improve the robustness of large-scale PIU networks.

Several limitations remain for future investigation. The present experiments validate the core PIU functions on a fabricated 4×4 MRR-based implementation, whereas larger-scale systems will require denser routing fabrics, improved thermal isolation and more automated calibration procedures. The hardware-aware inference simulations capture static channel-power mismatch and experimentally measured equalization capability, but do not yet include all nonidealities of a full photonic computing system, such as SOA noise, laser relative-intensity noise, detector noise, nonlinear gain saturation, dynamic traffic variation and long-term thermal drift across many cascaded nodes. Incorporating these effects into end-to-end photonic neural-network models will provide a more complete assessment of large-scale deployment.

Overall, the results show that interconnect-level equalization and stabilization are essential for preserving analog optical states in distributed photonic computing. The PIU therefore provides a practical route toward reconfigurable multi-core photonic systems in which optical signals can be routed between computing units while maintaining the power fidelity required for downstream computation.

4 Materials and methods

4.1 Chip design, fabrication and packaging

The chip was fabricated on a 180-nm silicon photonic platform. The MRRs have a radius of 10 μm, with a measured free spectral range of 8.715 nm and an average quality factor of $Q = 17{,}387$. The thermo-optic tuning efficiency was measured to be 0.26884 nm/mW. These parameters provide the wavelength selectivity and thermal tunability required for WDM routing, equalization and closed-loop resonance stabilization.

After fabrication, the chip was packaged for simultaneous optical and electrical access. Optical input and output signals were coupled between the fibers and the photonic chip through the packaged fiber interfaces. Electrical connections to the integrated heaters were provided through wire-bonded pads and external heater-driving probes.

4.2 Experimental setup

Closed-loop stabilization was implemented using a custom FPGA-based feedback controller built on an Altera/Intel Cyclone IV E FPGA device (EP4CE15F23C8). The monitored optical power was converted to an electrical signal and digitized by an 8-bit ADC. The FPGA compared the digitized power with the target value and updated the heater-driving voltage through an 8-bit DAC, thereby forming a real-time feedback loop for MRR resonance tracking and channel-power stabilization. Because the FPGA readout path was limited by the 8-bit ADC resolution, the steady-state optical-power fluctuations were additionally characterized using a Keithley 2602B source-measure unit, and these high-resolution measurements were used to extract the bit-equivalent control precision.

For temperature-dependent transmission measurements, the packaged PIU chip was mounted on a thermoelectric-cooler-controlled stage to apply controlled substrate-temperature perturbations. A 10-Gbaud NRZ pseudo-random binary sequence was generated by a Keysight M8045A high-speed pattern generator and applied to an Exail MXER-LN-20 lithium-niobate intensity modulator. The intensity-modulated optical signal was then transmitted through the PIU. After propagation through the selected MRR channel, the optical signal was detected by a photodetector followed by a transimpedance amplifier and recorded using a Keysight DCA-M N1092C digital communication analyzer for eye-diagram characterization. In the free-running measurements, the heater voltage was fixed after initial alignment. In the PID-locked measurements, the FPGA continuously updated the heater-driving voltage according to the monitored optical power to compensate resonance drift during temperature variation. The signal-to-noise ratio was extracted from the measured eye diagrams at each temperature point.

4.3 FPGA control algorithm

In Phase 1, a voltage–power lookup table (LUT) is constructed for the selected MRR by sweeping the heater-driving DAC code and recording the corresponding monitored optical power. The resulting LUT provides the accessible power range and the voltage–power mapping for the selected wavelength channel, but it is not used as a final locking bias because the MRR response is nonlinear, periodic, and susceptible to drift and discretization error.

In Phase 2, the FPGA performs LUT-assisted monotonic target acquisition using real-time ADC feedback. Rather than directly applying a LUT-inferred DAC code, the controller starts from the low-DAC side and incrementally increases the DAC code until the monitored power crosses the target level. The crossing criterion is port-dependent: for a drop-port response, the search stops when the ADC value rises above

the target, whereas for a through-port response, it stops when the ADC value falls below the target. After the target crossing is detected, the DAC code is backed off by a small predefined number of codes to reduce overshoot before entering the locking stage. This procedure uses the LUT to guide the target range while avoiding direct dependence on a pre-characterized bias point, thereby providing a more robust initialization under drift and device-state variation.

In Phase 3, sign-aware PID locking is activated to maintain the selected channel at the target power. Because the required heater-voltage correction direction depends on the selected resonance branch, the FPGA converts the raw power error into an effective error using an estimated branch polarity. The initial polarity is determined by the port configuration and is protected for a short interval after entering the locking state; subsequently, the polarity is updated from the measured signs of recent DAC-code and ADC changes. The PID loop further incorporates anti-windup protection through conditional integral accumulation, integral clamping, and DAC-range saturation. This sign-aware and anti-windup PID update enables stable closed-loop locking under branch-dependent MRR response, drift, and finite DAC range.

$$\begin{aligned} e_k &= \sigma_k \left(P_{\mathrm{tar}} - P_k\right), \\ u_{k+1} &= \mathrm{sat}_{[u_{\min}, u_{\max}]} \left[u_k + K_p e_k + I_k + K_d \left(e_k - e_{k-1}\right)\right]. \end{aligned} \tag{2}$$

where P_k and u_k are the measured power and DAC code at the k-th update, $\sigma_k \in \{+1, -1\}$ is the estimated polarity of the selected monotonic branch, and I_k is the bounded integral term. The integral term is conditionally accumulated and clamped to suppress windup, while $\mathrm{sat}_{[u_{\min}, u_{\max}]}(\cdot)$ limits the final DAC update to the available hardware range.

The sign-aware error preserves the correct feedback polarity on different resonance branches, while the anti-windup mechanism prevents excessive integral accumulation when the DAC code approaches its hardware limits.

4.4 Hardware-aware inference simulation

Hardware-aware inference simulations were performed using VGG-11 and ResNet-18 trained on the CIFAR-10 dataset. Both networks were trained under clean 32-bit floating-point (FP32) conditions, without injecting optical channel mismatch during training. Hardware impairments were introduced only during test-time inference by applying channel-dependent power gains to the input channels of each convolutional layer.

For the unequalized mismatch cases, the experimentally measured routed-channel powers were first converted into relative channel-loss classes by referencing the strongest routed channel. For each model, mismatch level, and Monte Carlo seed, static input-channel losses were sampled from these measurement-derived classes with a clipped dB-domain jitter and then kept fixed throughout inference. This procedure emulates a fixed hardware channel-mismatch realization caused by unequal routed-channel transmission. For the post-equalization case, the residual channel-power losses were uniformly bounded between 0 and 0.2 dB, consistent with the experimentally measured post-equalization variation.

To isolate channel-to-channel nonuniformity from common insertion loss, scalar layer-wise global-power compensation was applied by normalizing the sampled power gains by their layer-wise mean. All simulated conditions used the same baseline hardware-aware settings, including input Gaussian noise and low-bit input/weight quantization, with the detailed parameters provided in Supplementary Information. The layer-wise feature-map deviation was quantified using the relative L_2-norm difference between the impaired feature maps and the clean hardware-reference feature maps.

Data availability

The data supporting the findings of this study are available from the corresponding author upon reasonable request.

Code availability

The code used for the FPGA-based feedback control and hardware-aware inference simulations is available from the corresponding author upon reasonable request.

Acknowledgements

This publication is based upon work supported by the King Abdullah University of Science and Technology (KAUST) under Award Nos. RFS-TRG2024-6196, ORFS-CRG12-2024-6487, and RFS-sTRG2026-7454.

Author contributions

Z.L. and Z.Z. contributed equally to this work. Z.L. designed the photonic interconnect architecture, performed the device characterization, implemented the FPGA-assisted feedback-control experiments, carried out the hardware-aware inference simulations, analyzed the data and wrote the manuscript. Z.Z. contributed to the device design, experimental measurements and data analysis. X.O., W.L., X.W., H.Y., X.S. and W.H. contributed to device characterization, experimental setup development and technical discussions. Y.W. supervised the project, provided guidance on the research direction and revised the manuscript. All authors discussed the results and approved the manuscript.

Competing interests

The authors declare no competing interests.

Supplementary information

Supplementary information accompanies the manuscript on the Light: Science & Applications website.

References

[1] Shen, Y. *et al.* Deep learning with coherent nanophotonic circuits. *Nature Photonics* **11**, 441–446 (2017). URL https://doi.org/10.1038/nphoton.2017.93.

[2] Feldmann, J. *et al.* Parallel convolutional processing using an integrated photonic tensor core. *Nature* **589**, 52–58 (2021). URL https://doi.org/10.1038/s41586-020-03070-1.

[3] Xu, X. *et al.* 11 TOPS photonic convolutional accelerator for optical neural networks. *Nature* **589**, 44–51 (2021). URL https://doi.org/10.1038/s41586-020-03063-0.

[4] Shastri, B. J. *et al.* Photonics for artificial intelligence and neuromorphic computing. *Nature Photonics* **15**, 102–114 (2021).

[5] Zhou, H. *et al.* Photonic matrix multiplication lights up photonic accelerator and beyond. *Light: Science & Applications* **11**, 30 (2022).

[6] Sunny, F., Taheri, E., Nikdast, M. & Pasricha, S. *Silicon photonic 2.5d interposer networks for overcoming communication bottlenecks in scale-out machine learning hardware accelerators*, 1–4 (2024).

[7] Kannan, A., Jerger, N. E. & Loh, G. H. *Enabling interposer-based disintegration of multi-core processors*, MICRO-48, 546–558 (Association for Computing Machinery, New York, NY, USA, 2015).

[8] Kachris, C., Kanonakis, K. & Tomkos, I. Optical interconnection networks in data centers: Recent trends and future challenges. *IEEE Communications Magazine* **51**, 39–45 (2013).

[9] Cheng, Q., Rumley, S., Bahadori, M. & Bergman, K. Photonic switching in high performance datacenters [Invited]. *Optics Express* **26**, 16022 (2018).

[10] Zhang, W. *et al.* Silicon microring synapses enable photonic deep learning beyond 9-bit precision. *Optica* **9**, 579–584 (2022).

[11] Shafiee, A., Banerjee, S., Chakrabarty, K., Pasricha, S. & Nikdast, M. *LoCI: An Analysis of the Impact of Optical Loss and Crosstalk Noise in Integrated Silicon-Photonic Neural Networks*, GLSVLSI '22, 351–355 (Association for Computing Machinery, New York, NY, USA, 2022).

[12] Bogaerts, W. *et al.* Silicon microring resonators. *Laser & Photonics Reviews* **6**, 47–73 (2012).

[13] Tait, A. N. *et al.* Microring Weight Banks. *IEEE Journal of Selected Topics in Quantum Electronics* **22**, 312–325 (2016).

[14] Bai, B. *et al.* Microcomb-based integrated photonic processing unit. *Nature Communications* **14**, 66 (2023).

[15] Cheng, J. *et al.* Self-calibrating microring synapse with dual-wavelength synchronization. *Photonics Research* **11**, 347–356 (2023).

[16] Liu, X., Zhang, W., Cheng, J., Zhou, H. & Dong, J. Single-Monitor Calibration for Multiple Microring Synapses. *ACS Photonics* **11**, 2570–2577 (2024).

[17] Zhu, Q., Qiu, C., He, Y., Zhang, Y. & Su, Y. Self-homodyne wavelength locking of a silicon microring resonator. *Optics Express* **27**, 36625–36636 (2019).

[18] Echeverría, J. G. *et al.* Self-Calibrated Microring Weight Function for Neuromorphic Optical Computing. *Journal of Lightwave Technology* **43**, 602–610 (2025).

[19] Reck, M., Zeilinger, A., Bernstein, H. J. & Bertani, P. Experimental realization of any discrete unitary operator. *Physical review letters* **73**, 58 (1994).

[20] Clements, W. R., Humphreys, P. C., Metcalf, B. J., Kolthammer, W. S. & Walmsley, I. A. Optimal design for universal multiport interferometers. *Optica* **3**, 1460–1465 (2016).

[21] Spanke, R. Architectures for guided-wave optical space switching systems. *IEEE Communications Magazine* **25**, 42–48 (1987).

[22] Huang, Y. *et al.* Multi-Stage 8 × 8 Silicon Photonic Switch Based on Dual-Microring Switching Elements. *Journal of Lightwave Technology* **38**, 194–201 (2020).

[23] Lee, B. G. & Dupuis, N. Silicon photonic switch fabrics: Technology and architecture. *Journal of Lightwave Technology* **37**, 6–20 (2018).

[24] Cheng, Q. *et al.* Silicon photonic switch topologies and routing strategies for disaggregated data centers. *IEEE Journal of Selected Topics in Quantum Electronics* **26**, 1–10 (2019).

[25] Popovíc, M. A. *et al.* Multistage high-order microring-resonator add-drop filters. *Optics Letters* **31**, 2571–2573 (2006).

[26] Wu, Y. *et al.* Integrated photonic 3D tensor processing engine. *Light: Science & Applications* **15**, 154 (2026).

[27] Zhang, W., Wang, Y., Lederman, J. C., Shastri, B. J. & Prucnal, P. R. Compact, reconfigurable, and scalable photonic neurons by modulation-and-weighting microring resonators. *eLight* **6**, 6 (2026).

[28] Al-Fares, M., Loukissas, A. & Vahdat, A. *A scalable, commodity data center network architecture*, SIGCOMM '08, 63–74 (Association for Computing Machinery, New York, NY, USA, 2008). URL https://doi.org/10.1145/1402958.1402967.

[29] Sobhanan, A. *et al.* Semiconductor optical amplifiers: recent advances and applications. *Advances in Optics and Photonics* **14**, 571–651 (2022).

[30] Simonyan, K. & Zisserman, A. Very deep convolutional networks for large-scale image recognition. *arXiv preprint arXiv:1409.1556* (2014).

[31] He, K., Zhang, X., Ren, S. & Sun, J. *Deep residual learning for image recognition*, 770–778 (2016).

[32] Krizhevsky, A., Hinton, G. *et al.* Learning multiple layers of features from tiny images (2009).

[33] Verma, A. *et al. Large-scale cluster management at Google with Borg*, EuroSys '15, 1–17 (Association for Computing Machinery, New York, NY, USA, 2015).

[34] Benjamin, J. L., Gerard, T., Lavery, D., Bayvel, P. & Zervas, G. PULSE: Optical Circuit Switched Data Center Architecture Operating at Nanosecond Timescales. *Journal of Lightwave Technology* **38**, 4906–4921 (2020).

[35] Wang, Z., Ming, D., Wang, Y., Qiu, C. & Tan, M. Resolving the scalability challenge of wavelength locking for multiple micro-rings via pipelined time-division-multiplexing control. *Optics Express* **30**, 24984–24994 (2022).

[36] Shuangmei Xu, Khurgin, J., Vurgaftman, I. & Meyer, J. Reducing crosstalk and signal distortion in wavelength-division multiplexing by increasing carrier lifetimes in semiconductor optical amplifiers. *Journal of Lightwave Technology* **21**, 1474–1485 (2003).

[37] Mahajan, A., Christodoulopoulos, K., Martinez, R., Spadaro, S. & Munoz, R. Modeling EDFA Gain Ripple and Filter Penalties With Machine Learning for Accurate QoT Estimation. *Journal of Lightwave Technology* **38**, 2616–2629 (2020).

[38] Fan, Y. *et al.* Channel power equalization based on joint optimization of EDFA and ROADM configuration in open optical network. *Optics Express* **31**, 23229 (2023).

Supplementary Files

This is a list of supplementary files associated with this preprint. Click to download.

- SupplementaryInformation.pdf
- SupplementaryVideo.mp4