Flexible Error Modeling and Analysis of SiPho JTC Accelerators for CNNs

AARON YEN, University of California, Los Angeles, USA
BELAL JAHANNIA, University of Florida, USA
RUSSELL L.T. SCHWARTZ, University of Florida, USA
ANYU JIANG, University of Florida, USA
HANGBO YANG, University of Florida, USA
BAIBHAB CHATTERJEE, University of Florida, USA
VOLKER J. SORGER, University of Florida, USA
HAMED DALIR, University of Florida, USA
PUNEET GUPTA, University of California, Los Angeles, USA

This paper describes a system for modeling nonidealities in a silicon-photonic joint transform correlator (JTC) used as an accelerator for convolutional neural networks (CNNs). The system is used to determine which nonidealities have the greatest impact on end-to-end model accuracy, which helps determine what design choices can be made at the block and system level to improve performance. We create a digital twin of the on-chip 1D JTC using transfer functions fitted from simulations, then introduce a scalar α that mixes ideal and realized behaviors. We quantify the α-dependent model accuracy on CIFAR-10 and SNDR differences at both the joint-power spectrum (JPS) and the JTC output. The framework allows quantification of block- and system-level sensitivities, informing design choices by isolating the accuracy bottlenecks within the system.

CCS Concepts: • **Hardware → Analysis and design of emerging devices and systems**; • **Computing methodologies**;

Additional Key Words and Phrases: Silicon photonics, Joint transform correlator, Convolutional neural networks, Accelerator architectures, Digital twins

ACM Reference Format:
Aaron Yen, Belal Jahannia, Russell L.T. Schwartz, Anyu Jiang, Hangbo Yang, Baibhab Chatterjee, Volker J. Sorger, Hamed Dalir, and Puneet Gupta. 2026. Flexible Error Modeling and Analysis of SiPho JTC Accelerators for CNNs. *ACM Trans. Des. Autom. Electron. Syst.* 1, 1 (June 2026), 21 pages. https://doi.org/10.1145/nnnnnnn.nnnnnnn

1 Introduction

The classical Joint Transform Correlator (JTC) utilizes a property of Fourier optics wherein a spatial signal placed at the front focal plane of an optical lens will yield its Fourier transform (FT) at the back focal plane [12]. We can leverage this property to convolve two spatial signals by placing them together at the front focal plane of an optical lens. An array of optical modulators modulates the signals from an initial light source. The output at the back focal plane allows us to calculate the joint power spectrum (JPS). Performing the inverse Fourier transform (IFT) on the JPS yields the correlation output. Spatially, this result includes both auto-correlation and cross-correlation terms. For real-valued inputs, the

Authors' Contact Information: Aaron Yen, aaronhyen@ucla.edu, University of California, Los Angeles, Los Angeles, California, USA; Belal Jahannia, University of Florida, Gainesville, Florida, USA; Russell L.T. Schwartz, University of Florida, Gainesville, Florida, USA; Anyu Jiang, University of Florida, Gainesville, Florida, USA; Hangbo Yang, University of Florida, Gainesville, Florida, USA; Baibhab Chatterjee, University of Florida, Gainesville, Florida, USA; Volker J. Sorger, University of Florida, Gainesville, Florida, USA; Hamed Dalir, University of Florida, Gainesville, Florida, USA; Puneet Gupta, puneetg@ucla.edu, University of California, Los Angeles, Los Angeles, California, USA.

2026. Manuscript submitted to ACM

cross-correlation term is a mirrored version of the convolution of the two input spatial signals. Thus, we can use JTC to perform the convolution operation commonly used in machine learning (ML) applications. Recent developments in silicon photonics (SiPho) and analog mixed-signal (AMS) circuits for communications applications, such as high-speed microring modulators (MRMs), digital-to-analog converters (DACs), analog-to-digital converters (ADCs), and driver circuits, enable integration of 1D on-chip lenses, AMS circuits, and SiPho circuits on a 45nm monolithic process [30] to fabricate a tightly integrated replication of the bulk optics flow for a JTC on a single chip, enabling high-throughput low-power convolutional neural network (CNN) accelerators [20]. This work presents a framework for analyzing photonic CNN accelerators by cataloging nonideal behaviors at each point in an on-chip JTC.[1]

This work analyzes a 1D JTC-based photonic CNN accelerator as described in [20]. We develop a digital twin emulating the ideal and behavioral characteristics of each element in the system. We use a weighted average of the ideal and realized responses as a sensitivity knob to study how moving toward ideal behavior (e.g., via circuit/process improvements or calibration) affects system-level performance.

2 Background and Related Work

2.1 Fourier-Domain Convolution and Optical Correlators

The primitive for CNNs is the spatial convolution operation, where a kernel of dimensions $k \times k$ is applied at each position of an $N \times N$ matrix. While traditional spatial convolution operation has a worst-case time complexity of $O(N^2k^2)$, convolution can also be performed in the spatial Fourier domain by transforming the input matrix into the frequency domain, applying the matrix as an element-wise multiplication, then taking the inverse Fourier transform to return to the spatial domain. With the Fast Fourier Transform (FFT) algorithm, the worst-case time complexity is $O(N \log(N))$ for a 1D signal and $O(N^2 \log(N))$ for a 2D image. From Fourier optics, we know that when coherent light passes through an optical lens, the lens will project the Fourier transform of the signal at the input focal plane on the output focal plane. By using two lenses and a spatial light modulator (SLM) in the Fourier plane, the time complexity of such a convolution can be reduced to $O(N^2)$, the complexity of modulating signals at the input and Fourier planes. This configuration is known as the 4F correlator for the four focal length distances between the input and output planes.

2.2 Joint Transform Correlator (JTC) Operation

The accelerator class analyzed in this work is an optical correlator variant known as the Joint Transform Correlator. The major limitations of the 4F correlator are strict alignment requirements to ensure the filter is located correctly in the frequency domain and the need for both amplitude and phase modulation of the filter.

In the JTC, the input signal $s(x)$ and kernel $k(x)$ are amplitude-only combined and passed through a lens to generate their combined Fourier transform in the optical domain, represented as

$$E_{\text{in}}(x) = s\left(x + \frac{\Delta}{2}\right) + k\left(x - \frac{\Delta}{2}\right) \tag{1}$$

where $s(x)$ and $k(x)$ are separated by Δ. $E_{\text{in}}(x)$ is passed through an optical lens performing the Fourier transform

$$E_{\text{F}}(u) = S(u)\, e^{+j2\pi u\, \Delta/2} + K(u)\, e^{-j2\pi u\, \Delta/2} \tag{2}$$

[1]Corresponding code is available at https://github.com/nanocad-lab/onn-flex

which is then captured by an array of square-law detectors to obtain $I_F(u)$, another amplitude-only signal.

$$\begin{aligned} I_F(u) &= |S(u)|^2 + |K(u)|^2 \\ &+ S(u)K^*(u)e^{j2\pi\Delta u} \\ &+ S^*(u)K(u)e^{-j2\pi\Delta u} \end{aligned} \tag{3}$$

This intensity pattern $I_F(u)$ is passed through a second lens, producing an output $E_{JTC}(x)$ that is a composite of the auto-correlations and cross-correlations of $s(x)$ and $k(x)$.

$$\begin{aligned} E_{JTC}(x) =& \mathcal{F}^{-1}\{I_F(u)\}(x) \\ &= s(x) * s(-x) + k(-x) * k(x) \\ &+ s(x - \Delta) * k(-x) \\ &+ s(-x) * k(x + \Delta) \end{aligned} \tag{4}$$

As the JTC only requires modulation of amplitude or intensity, it does not incur the additional complexity of phase modulation. Additionally, as the signal and kernel are passed through the lens together, the alignment requirement is more lenient compared to 4F correlator systems.

2.3 On-Chip Silicon-Photonic JTC (pJTC) Architecture

In this work, we analyze the performance of a proposed silicon-photonic JTC accelerator that leverages high-speed optical-transceiver building blocks to shrink the JTC system to a chip-scale solution. As shown in Figure 1, high-speed DACs and differential drivers control an array of MRMs feeding a slot-based Fresnel lens [40] that performs a Fourier transform. The output is captured by photodetectors (PDs) and then converted by matched transimpedance amplifiers (TIAs) and ADCs. This enables a single-lens JTC implementation; however, it introduces quantization noise between the first and second Fourier transform operations. The resulting link is constrained by optical power budgeting across the modulator array, diffraction efficiency and phase errors of the integrated lens, and the dynamic range available after electronic amplification and digitization. These constraints define an operating window where photonic shot noise, MRM extinction ratios, and ADC quantization jointly shape the usable JPS fidelity, directly informing the nonideality models we integrate.

2.4 Related Work on Optical Correlators and Photonic Accelerators

Previous works on JTC error minimization included binarization of the JPS [13, 15, 16], phase shifting and encoding [24, 34], and fringe-adjusted filters [1, 7, 38]. For the CNN accelerator use case, binarization of the JPS is not tenable as the JTC is used for convolution, not object detection. Binarizing the convolution operation can be effective for specific networks and datasets, but it is not a universal replacement for convolution. Similarly, phase-shifting and encoding were developed for detecting correlation peaks and are not a stand-in for convolution. However, these methods may be more beneficial at a model level, where such peaking can function similarly to an activation function.

Most silicon photonic neural network accelerators (SPNNAs) today are focused on performing matrix-vector multiplications [28]. Previous works analyzing error sources in such accelerators have identified error sources such as thermal crosstalk or process variations affecting waveguide lengths [2, 3, 9, 33]. Thermal crosstalk is a critical concern for mesh-based SPNNAs that utilize heaters to control MZIs or MRMs [22, 36]. However, our pJTC system uses P-N

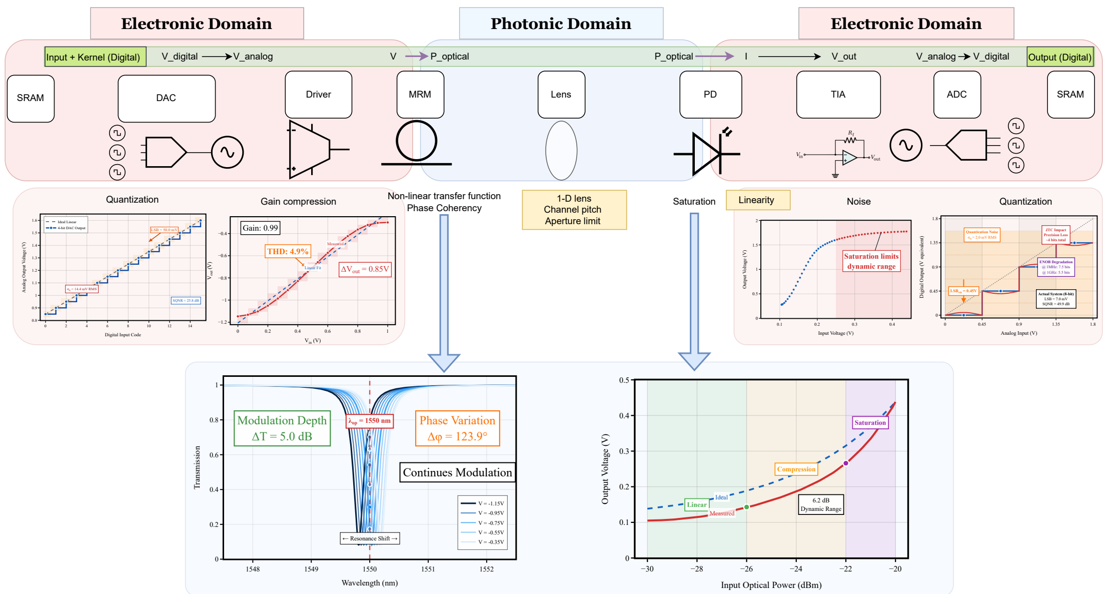

Fig. 1. Data flow of one partial forward pass through the physical components of the on-chip JTC

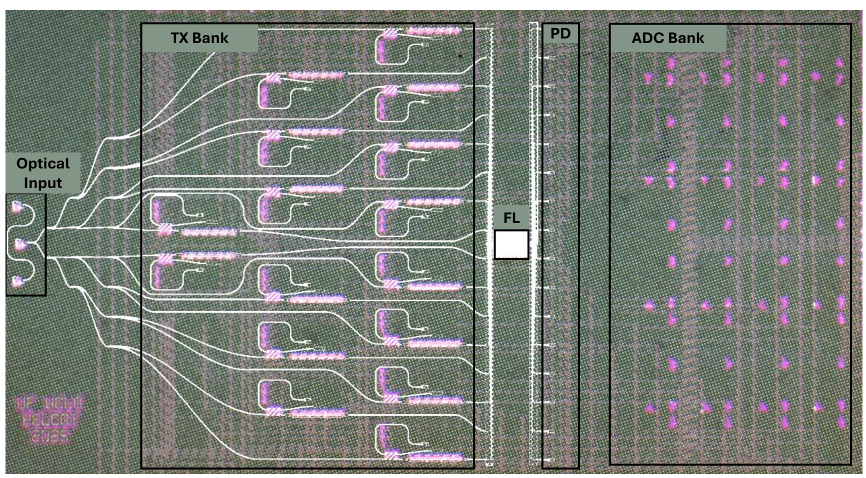

Fig. 2. Annotated die image of the pJTC test-chip design (fabricated in GlobalFoundries' 45SPCLO (45CLO) process), highlighting the optical input, transmitter bank, Fresnel lens (FL), photodetector (PD) region, and ADC bank.

junctions to control MRM output, allowing any thermal gradient caused by heaters to be corrected over a much longer timescale without performance impact. Similarly, as the majority of high-speed optical datacenter interconnects (DCI) utilize intensity modulation direct detection (IMDD) for lower complexity and cost [4], analyses of MRM performance for high-speed links are primarily focused on its qualities as an amplitude-only modulator. However, the field at the optical lens input can accumulate phase errors from waveguide-length variations (process or thermal) in addition to the

MRMs' parasitic phase modulation. Both effects perturb the FT output, since the ideal JTC input signal contains no phase modulation.

As many of the components of the on-chip JTC were initially developed for high-speed communications, many of the limitations are shared between the communication and pJTC computation applications of these devices. For example, MRM thermal resonance locking represents an energy overhead [23] in both systems. While existing works analyzing link budget [37] for optical DCI typically budget for insertion losses and crosstalk not seen in an on-chip system, the elements that contribute to performance degradation, such as shot and thermal noise [25, 31], are consistent between the two systems.

3 Digital-Twin Framework and Evaluation Methodology

3.1 Transfer-Function Modeling and α-Mixing

Using GlobalFoundries' 45SPCLO PDK device/circuit models (characterized and validated against fabricated-chip measurements by the platform/PDK developers [5, 27, 30]) and standard electronic design automation (EDA) and circuit/photonic simulation tools (e.g., Cadence Virtuoso), we generate circuit/photonic simulation sweeps (with post-layout parasitic extraction where applicable) for each element in the on-chip JTC shown in Figure 1. In this paper, all reported block transfer functions are derived from these simulation sweeps. To capture each component's impact, we model each element as a steady-state (memoryless) transfer function between its inputs and outputs. This allows us to incorporate component-level simulations while maintaining a differentiable representation for ML training. The input/output sweep data are fit with an n-th order polynomial where n is chosen to minimize the Akaike Information Criterion (AIC). The pJTC test chip integrates a 16-channel Fresnel lens accommodating an 8-channel input and 8-channel kernel and sampling an 8-channel output. The system operates using an external 1310 nm laser as the optical power source and is paired with an FPGA to handle digital processing and all elements not shown in Figure 1. We refer readers to [40] for experimental characterization data on the pJTC platform. Since our block models are static, we assume the system is clocked such that bandwidth/transient effects are negligible and intersymbol interference (ISI) is insignificant over the short on-chip channel.

To model the impact of each component in the JTC system, we create a digital twin by collecting detailed input/output responses for each element from PDK-based simulation sweeps. When additional hardware measurements are available, the same methodology can directly fit transfer functions to measured I/O sweeps without changing the overall framework.

Table 1. Summary of PDK-based simulation models used in the digital twin. H_{fit}^{n} is an n-th order polynomial fit to the corresponding input/output sweep data and R^2 is computed on the sweep data.

Block	**I/O range**	H_{ideal}	n	R^2	**Data source**
Driver	$[0, 1] \to [-1.147, -0.3]$ V	Linear	3	0.999979	Post-layout circuit simulation
MRM amplitude	$[-1.147, -0.3]$ V $\to [1.20, 2.14] \times 10^{-3}$ W$^{1/2}$	Linear	2	0.999614	Device/circuit simulation
MRM phase	$[-1.147, -0.3]$ V$\to [-0.593, 0.593]$ rad	0 (ideal amplitude-only)	2	0.999505	Device/circuit simulation
PD	$[1, 10]$ μW$\to [0.105, 0.438]$ V	Linear	2	0.999763	Circuit simulation
TIA	$[0.105, 0.438]$ V$\to [0, 1]$ (norm.)	Linear	6	0.999726	Post-layout circuit simulation

$$H_{\alpha}(x) = \alpha \cdot H_{\text{fit}}^{n}(x) + (1 - \alpha) \cdot H_{ideal}(x) \tag{5}$$

To analyze the impact of nonideal behavior in each component, we use $H_{\alpha}(x)$ as described in Equation 5, where the ideal transfer function $H_{ideal}(x)$ is linearly combined with the n-th order fitted polynomial transfer function $H_{\text{fit}}^{n}(x)$ to

Table 2. Notation for the JTC input-field layout parameterization.

Symbol	Definition
M	Input signal length (number of elements)
N	Kernel signal length (number of elements)
L	Total lens field size (number of channels)
Δ	Center-to-center separation between input and kernel
P_T	Top padding: unused channels above the input signal
P_M	Middle padding: unused channels between input and kernel
P_B	Bottom padding: unused channels below the kernel

form a weighted sum. Then, α is varied from 0 to 1 to perform design-space exploration (DSE) and sensitivity analysis. The transfer functions are applied in succession to simulate the behavior of an on-chip JTC, and we implement the same component chain as a PyTorch model to measure its utility in a machine learning context. The benefit of such a representation is that by characterizing the individual blocks, we can use lightweight mathematical models to capture the behavior of each block rather than constantly running computationally-intensive Spice or photonic simulations. At the scale of ML models, this significantly improves the speed at which analysis can be performed. Additionally, backpropagation can be applied to the mathematical models, such that network weights can be trained to compensate for some nonidealities of the system. The term α is used as a continuous interpolation for sensitivity analysis and does not necessarily correspond to a specific process knob or hardware-improvement trajectory; however, when comparing candidate design changes (e.g., improving MRM linearity), sweeping α can serve as a coarse proxy for incremental improvement along that development direction.

3.2 ML Workload and Training Setup

To benchmark the accuracy of our proposed accelerator, we evaluate its performance on supervised image classification of the CIFAR-10 dataset, reporting top-1 accuracy on the official test set. Convolution layers are realized using the pJTC pipeline with each convolutional layer consisting of stitched 1D convolution segments with size limitations dictated by the maximum input, kernel, and output length supported by the pJTC hardware. As optical intensities are nonnegative, the signed filter for each layer is represented by two unsigned filters whose outputs are subtracted to realize a differential filter pair.

The simulated model is trained using AdamW with cosine annealing and an initial learning rate of 10^{-3} for 20 epochs with a batch size of 128. We evaluate VGG-11 on CIFAR-10 to demonstrate the high accuracy capability of the pJTC accelerator. Table 3 shows that the ideally modeled optical pipeline performs within 1.58 percentage points of the fp32 digital baseline, while the 4-bit DAC and 6-bit ADC quantized optical model reduces accuracy to 82.92%, comparable to the 4-bit weight and activation quantized digital reference (80.9%).

Table 3. CIFAR-10 accuracy on VGG-11 under digital and optical execution. Δ values are in percentage points relative to the digital baseline; small discrepancies may arise from rounding.

Nonideality / Setting	Accuracy	Δ (%)
Baseline (Digital)	91.9%	—
Digital (4-bit)	80.9%	-11.0
Optical (Ideal)	90.3%	-1.58
Optical (Quantized)	82.92%	-8.98

While VGG-11 can achieve a higher accuracy score, the large design space requires many repeated training runs across nonideality settings and training configurations. Therefore, we use a shallower CNN for the remainder of the experimentation to keep runtime and memory usage manageable. This network begins with a 16-channel photonic convolution, batch normalization, ReLU, and 2×2 max pooling, followed by a 32-channel photonic convolution with the same sequence. This is followed by a configurable stack of zero or more identical blocks, each containing a 32-channel photonic convolution, batch normalization, and ReLU. Finally, the classifier head consists of 2×2 max pooling, flatten, and two fully connected layers of size 256 and 10. Kernel length is fixed at 8 to match the correlator.

3.3 Evaluation Metrics

For each component of the JTC, we analyze the impact of its nonidealities at three junctures:

(1) Signal-to-noise and distortion ratio (SNDR) at the ADC output of the first Fourier pass, computed against an ideal Fourier transform (JPS).
(2) SNDR at the ADC output of the second Fourier pass, computed against an ideal 1D convolution.
(3) Top-1 accuracy of an ML model trained on an ideal JTC system when executed on a nonideal system.

For SNDR calculations, the outputs of the JTC system and the ideal output are normalized to the [0, 1] range, computed per sample, and averaged over 100,000 random inputs.

For accuracy baselines in the design-space sweeps (shallow network), we also evaluate an all-digital reference using the same network topology, but with photonic convolution blocks replaced by standard digital convolution. This FP32 baseline reaches 67.5% accuracy. To isolate quantization effects, we insert the same quantize/dequantize operators used in our optical pipeline at the corresponding points (4-bit inputs and weights and 6-bit outputs), which yields 60.13% accuracy. To approximate the single-lens pJTC's Fourier-plane sampling, we additionally apply 4-bit quantization in the Fourier domain, reducing the digital model to 46.66% accuracy, comparable to the ideal optical model.

For SNDR comparisons, we expect a theoretical maximum achievable SNDR of 37.88 dB, equivalent to 6 bits ENOB (Effective number of bits), such that the difference between ideal and simulated outputs is less than half of the LSB of the ADC.

4 Component-Level Nonideality Models and Sensitivity

4.1 Modulator Stage (MRM): Amplitude Linearity and Parasitic Phase Modulation

The ideal modulator for a JTC is a linear, amplitude-only spatial light modulator (SLM). In practice, high-speed, low-power modulators such as MRMs are not ideal and thus require nonlinear input schemes to achieve linearized output amplitude [8]. Another characteristic of MRMs is that they modulate phase in addition to magnitude [41]. In high-speed optical communications, this manifests as a frequency "chirp".

$$E_{postlens} \propto \int A(x)e^{i\phi(x)}e^{-i\frac{2\pi}{\lambda f}xu}dx \tag{6}$$

In JTC applications, parasitic phase modulation presents as a phase distortion at the input of the on-chip lens. Equation 6 represents the field at the output plane of the on-chip lens where an optical signal with amplitude $A(x)$ and phase $\phi(x)$ is applied at the input plane, where x is the coordinate in the input plane, u is the coordinate in the output plane, λ is the modulated wavelength, and f is the focal length of the lens. An ideal SLM will produce no phase modulation, *i.e.*, $\phi(x) = 0$, whereas MRMs, with input-dependent $\phi(x)$, will distort the output.

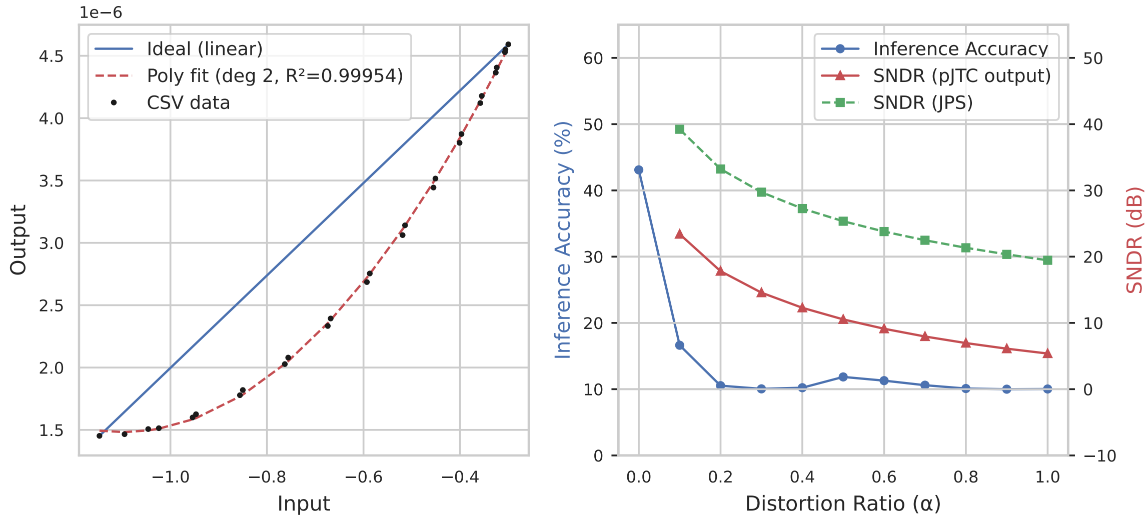

Fig. 3. MRM amplitude response: ideal vs realized transfer function (left) and impact on model accuracy and convolution SNDR (right).

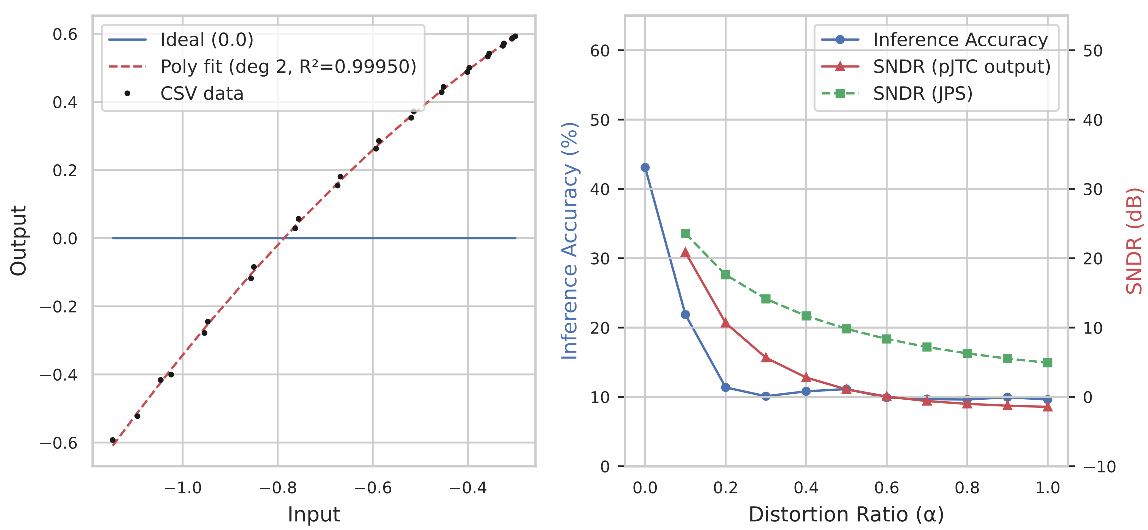

Fig. 4. MRM phase response: ideal vs realized transfer function (left) and impact on model accuracy and convolution SNDR (right).

To test the impact of varying levels of nonideality, we use α to represent the proportion of the resultant transfer function $H(x)$ that is from an n-th order polynomial fitted to simulation data representing the input-output relationship for a given component of the JTC engine. A possible solution to unwanted phase modulation is the use of microring-assisted Mach-Zehnder modulators and other push-pull dual MRM configurations that have been demonstrated to perform amplitude-only modulation for high-speed communications [6, 11, 19, 39, 43, 44] up to 124 GBaud [10]. These configurations minimize chirp by combining the output of two MRMs driven with opposite values, nulling their respective phase modulations.

A second limitation of MRMs is that the output is not linear with the voltage applied, but rather is dependent on the quality factor of the resonator to determine the steepness of its voltage curve. A possible solution is to design a ring with a lower Q factor and operate in the linear slope regime. Reducing Q would increase the range where the detuning slope is linear (in WL/Hz), translating to a larger V range. It is also possible to limit the electrical input to the MRM to operate within the quasi-linear region of the modulator.

This can be resolved by using a higher resolution DAC and applying a pre-distortion look-up table that corrects for MRM nonlinearity. It is also important to note that MRMs exhibit asymmetric rise and fall times due to carrier recombination times. At higher symbol rates, the transmitter may need to implement an asymmetric, pattern-dependent pre-emphasis filter.

4.2 Data Converter (DAC/ADC) Linearity and Quantization

The nonideal transfer characteristics of the onboard DACs and ADCs are modeled using a similar transfer-function approach. The ADC and DAC are modeled using a quantize-dequantize scheme with a straight-through estimator in the backward pass, with the fitted nonlinear transfer function applied before and after the quantization, respectively. For our steady-state analysis, we use the standard DNL and INL metrics. DNL measures the step-to-step deviation, and INL measures the overall deviation of the converter. Aside from nonidealities for these components not unique to the JTC, their input/output must be well integrated with the rest of the JTC system to maximize the available dynamic range. In our model, we use the transition levels for each code to build a nonuniform quantizer that can be applied at analog/digital transitions.

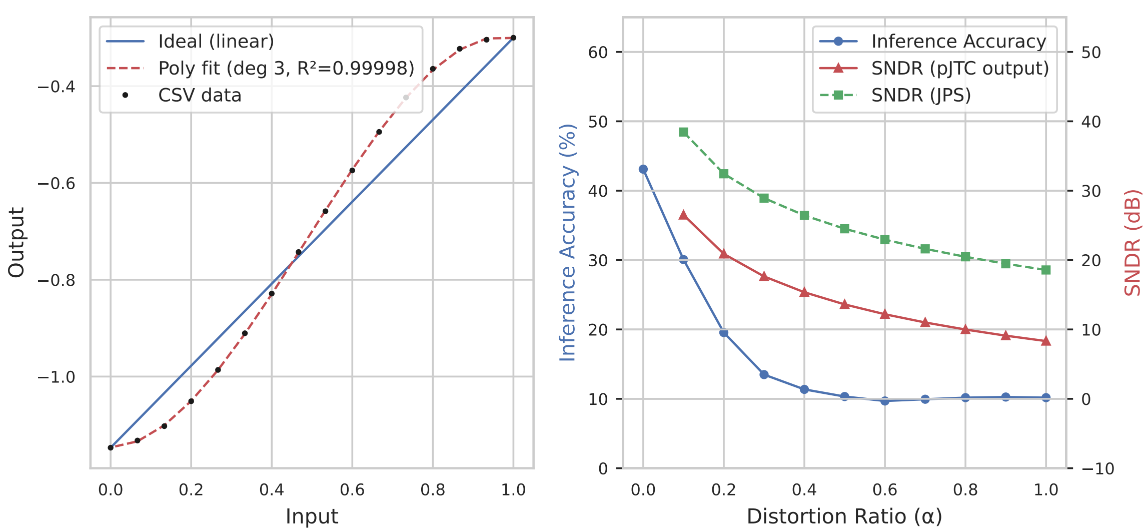

Fig. 5. Driver response: ideal vs realized transfer function (left) and impact of nonlinearity on model accuracy and convolution SNDR (right).

4.3 Electrical Driver Linearity

In the JTC system, the driver functions as an interface between the DAC output and MRM input, providing impedance matching, current drive capacity, and scaling. In our case, the driver's adjustable bias current allows us to tune the output voltage range to match the MRM's input voltage range. The requirements and limitations for an MRM driver for computation and communications systems are similar, and as such are not discussed in detail. It can be seen that despite the small (1.2%) deviation of the driver transfer function from the ideal relationship at $\alpha = 0.1$, Figure 5 shows a significant impact on model accuracy with a drop from 43.1% to 30.1%. This is largely due to distribution shift rather than information loss, as the difference in transfer function acts similarly to an input-dependent gain error, shifting activation statistics away from trained values. By re-estimating only batch-normalization statistics, we can recover accuracy from 30.1% to 41.1% at $\alpha = 0.1$, close to the 43.1% ideal baseline.

4.4 Photodetector and TIA: Linearity and Saturation

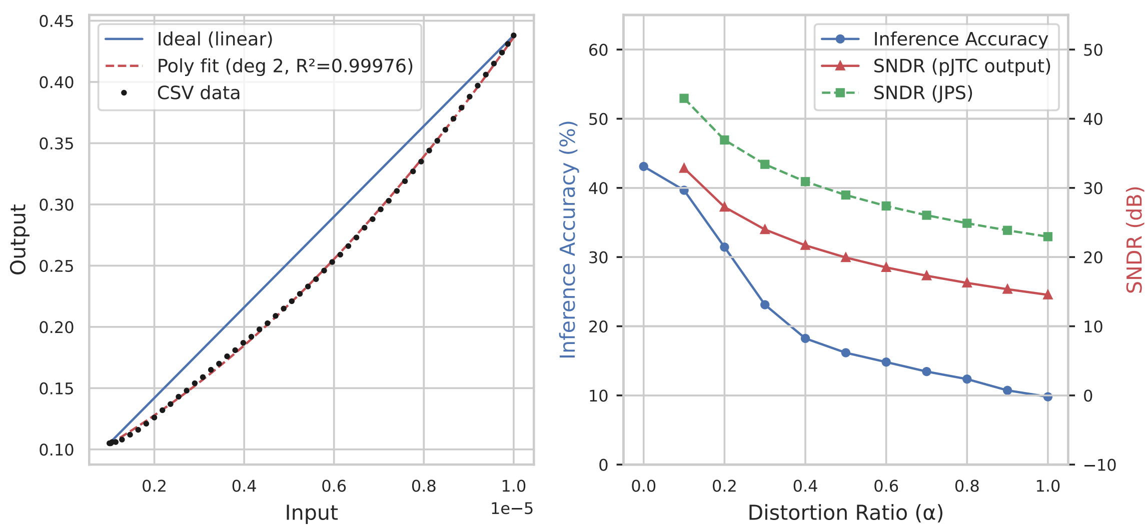

Fig. 6. PD response: ideal vs realized transfer function (left) and impact on model accuracy and convolution SNDR (right).

It is also critical that the system utilizes the optimal operating region of the PD and TIA. The laser power and the MRMs' optical modulation amplitude (OMA) control the PD input optical power range. A tradeoff exists between increasing the laser power to improve PD detection precision and remaining within the PD's linear and saturation-free input range. Similarly, the TIA should be designed and configured to accept the PD output within its linear range. It should also scale the output to the optimal level for ADC operation, as optimal clipping can improve quantized performance by as much as 1 bit equivalent [32].

Figure 6 shows that the ideal receiver transfer function has a linear relationship between detected optical power and voltage. However, we can see from Figures 7 and 8 that TIA saturation limits the linearity of the PD/TIA response.

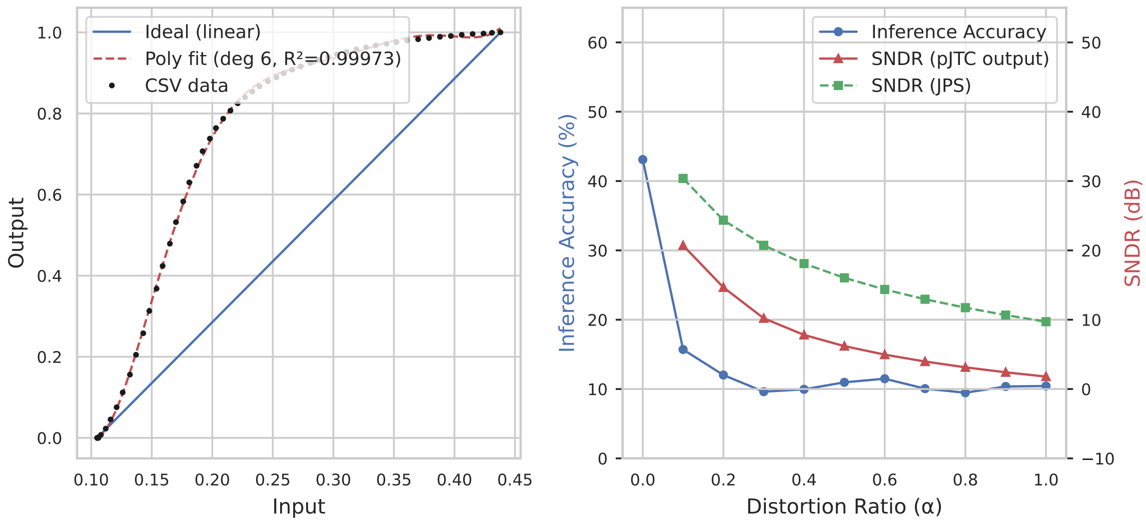

Fig. 7. TIA response: ideal vs realized transfer function (left) and impact on model accuracy and convolution SNDR (right).

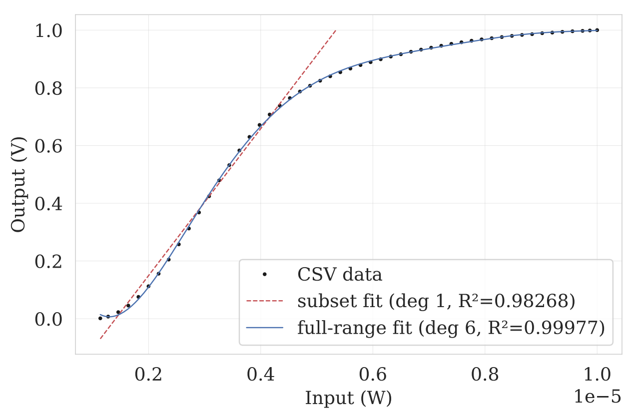

Fig. 8. Composite PD/TIA response: a linear fit to the pre-saturation region against the full-range sixth-order fit, isolating the deviation caused by TIA saturation.

5 System-Level Architectural Effects

5.1 Integrated Lens Model

For the on-chip JTC, a 1D Fresnel lens is used to realize the Fourier transform. This introduces optical aberrations, including discretization from Fresnel-lens faceting and manufacturing variations from process imperfections. Such imperfections can be captured by applying a distortion operator on the Fourier-plane field. Here, we assume that the pixel pitches of the input and output waveguides to the lens are equivalent and thus model the lens as an ideal FFT. We focus on the design space decision for the size of the optical field in points, or the number of input and output channels available on each focal plane of the lens. Increasing the size of the optical field allows for separation of the input signal and the convolution kernel, which is necessary to separate the auto-correlation and cross-correlation terms at the JTC output [35]. In this paper, the geometric field-size and placement sweeps in Sec. 5.4 and 5.5 assume an ideal FFT lens. As a result, those conclusions should be interpreted as *non-lens-limited*, isolating geometric effects and interactions with other modeled nonidealities.

To quantify lens nonideality in a compact, system-level form, we represent the realized lens as a discrete linear operator Q mapping input waveguide channels to output channels (after normalizing out overall loss) [12]. For the sensitivity study in this paper, we use a representative Q obtained from a scalar-diffraction simulation of an integrated 1D Fresnel lens [12]. Let F denote the ideal (unitary) discrete Fourier matrix (so F^{-1} is its inverse/conjugate-transpose). After removing a global phase, we define the unitary deviation $U = QF^{-1}$ and parameterize it as $U = \exp(i\Theta)$ with Hermitian generator Θ [14].

$$U = QF^{-1}, \qquad U = \exp(i\Theta) \tag{7}$$

We then fit a reduced-order generator using a 1D Legendre-induced basis by conjugating a diagonal Legendre phase profile by the ideal unitary DFT pair (F and F^{-1}) [12, 17, 21].

$$\Theta \approx \Theta_{\text{fit}} = \sum_{k=1}^{K} a_k B_k, \qquad B_k = F\,\text{diag}(\tilde{P}_k)\,F^{-1} \tag{8}$$

This construction corresponds to applying a diagonal Legendre phase profile across the lens channels and mapping it through the ideal Fourier operator. Here $\tilde{P}_k$ is the sampled, piston-removed, normalized k-th Legendre mode on $\rho \in [-1, 1]$. Figure 9 shows the system-level impact of the 2nd-order lens model across varying values of distortion strength α.

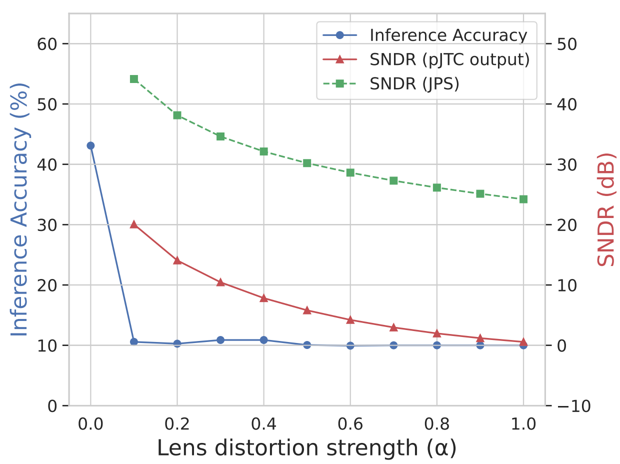

Fig. 9. Impact of lens distortion strength (α) on inference accuracy and SNDR at the JPS and pJTC output.

5.2 Process Variation: Splitter Imbalance

Process variations are a critical concern for pJTC systems. [9] shows that variations in line edge roughness in silicon photonics processes can result in significant imbalances in a 50-50 optical splitter. This can be somewhat improved with careful design [9], but still has a large potential impact on the system's performance [42]. As our system is sensitive to both phase misalignment and amplitude imbalance, we can expect that process variations will have a significant impact on the system's performance.

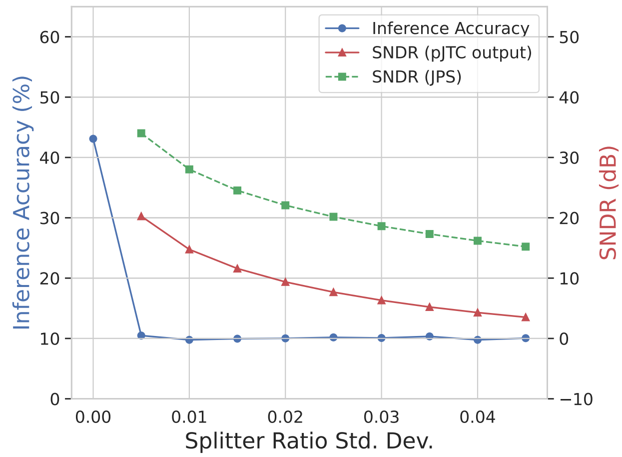

Fig. 10. pJTC performance across different splitter imbalance distributions.

In our design, we use a series of 1⟶2 splitters in a balanced binary tree configuration to create an array of waveguides for modulation. We model the possible variations in the MRM input optical amplitude by applying batched Monte-Carlo draws for each of the splitters in the tree. Each splitter's ratio is sampled from $\mathcal{N}(0.50,\ \sigma^2)$. This simulates inference on an array of pJTC cores, each having slightly different imbalances. Figure 10 shows how widening splitter imbalance distributions impacts pJTC inference accuracy.

It is feasible to define a procedure to calibrate individual array amplitudes by modifying the global output amplitude at the DAC or driver output, or through the use of variable optical attenuators. A set of patterns selected to remove lens output channel variation impacts could be used to isolate the amplitude variations of individual channels. Similarly, one could select patterns to calibrate for variations at the output of the lens prior to the ADC.

5.3 Stochastic Noise: Laser Amplitude Noise and PD Input Noise

To address additional noise sources beyond deterministic transfer-function distortions, we add two effective stochastic noise models that operate on optical power and are sampled independently for the two optical passes (JPS and output-plane measurement). We parameterize laser source amplitude noise by a relative RMS intensity level σ_{rin} (reported in dB as $20\log_{10}\sigma_{\mathrm{rin}}$), and we parameterize receiver noise by an input-referred PD optical-power noise standard deviation σ_{pd} (W RMS). First, we model laser amplitude noise as a per-shot multiplicative perturbation on optical power. For each optical pass and each shot, we sample $\varepsilon \sim \mathcal{N}(0,1)$ and define a log-normal scale factor $s = \exp\left(\sigma_{\log}\varepsilon - \frac{1}{2}\sigma_{\log}^2\right)$, where $\sigma_{\log}$ is the standard deviation of $\ln s$ and the $-\frac{1}{2}\sigma_{\log}^2$ term enforces $\mathbb{E}[s] = 1$. We define the target relative RMS intensity noise as $\sigma_{\mathrm{rin}} = \sigma_I/I$, where I is the nominal (mean) optical power for that shot and σ_I is the standard deviation of the per-shot intensity fluctuation (equivalently, since $\mathbb{E}[s] = 1$, σ_{rin} is the square root of the *variance* of s). For the unit-mean log-normal model above, this yields $\sigma_{\log} = \sqrt{\ln(1+\sigma_{\mathrm{rin}}^2)}$. The same s is broadcast across channels within a shot and is applied prior to splitter imbalance, so common-mode source noise becomes channel-uneven after the nonideal splitter tree. σ_{rin} is varied as an effective integrated amplitude-noise metric (not a RIN PSD in dB/Hz). For a flat, one-sided RIN power spectral density of the normalized power fluctuation ($\Delta P/P$) with linear units RIN_{lin} (1/Hz) over an effective noise-equivalent bandwidth B_{NE} (Hz), the relative RMS satisfies $\sigma^2_{\text{rin}} \approx \text{RIN}_{\text{lin}} B_{\text{NE}}$.

Second, we model photodetector noise as an input-referred additive Gaussian perturbation on the PD input optical power, sampled i.i.d. across tensor elements and clamped to the modeled PD power range. We sweep σ_{rin} from -80 to -30 dB (0.01% to 3.16% RMS) and σ_{pd} from 10^{-10} to 10^{-6} W RMS to span receiver-relevant regimes. For example, [18] reports a DFB laser with measured RIN < -150 dBc/Hz and an on-chip PD with responsivity ≈ 1.02 A/W and O/E bandwidth > 20 GHz, while [26] estimates an input-referred receiver noise density of 14 pA/$\sqrt{\text{Hz}}$. Using $B \approx 10$ GHz gives $\sigma_{\text{rin}} \approx 3 \times 10^{-3}$ (i.e., ≈ -50 dB, or 0.3% RMS) for RIN $= -150$ dBc/Hz, and shot-noise-limited detection gives $\sigma_P \approx \sqrt{2qPB/R}$, which evaluates to $\sigma_P \sim 10^{-7}$ W for $P \sim 10^{-6}$–10^{-5} W, $R \sim 1$ A/W, and $B \sim 10$–20 GHz [18, 29].

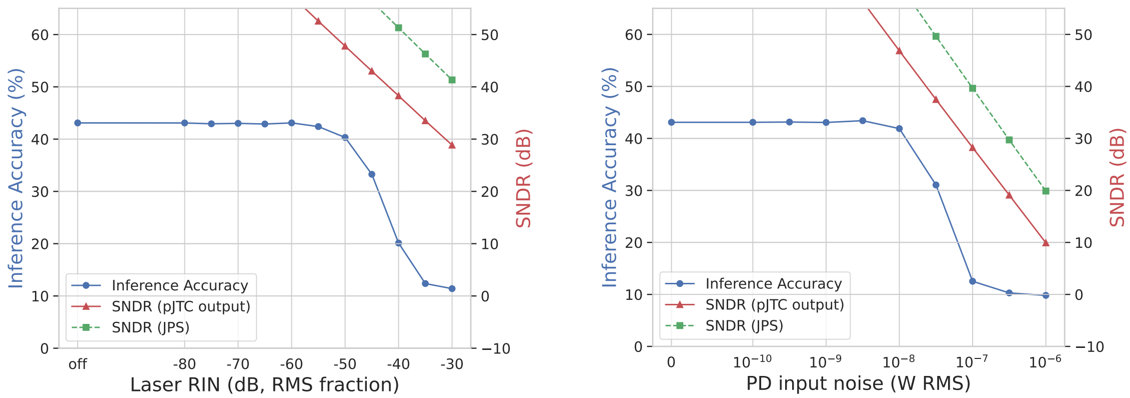

(a) Laser amplitude-noise sweep.

(b) PD input-referred noise sweep.

Fig. 11. Impact of effective stochastic noise sources on inference accuracy and SNDR at the JPS and pJTC output.

Figure 11a shows that laser amplitude noise has a limited impact on accuracy for σ_{rin} below -50 dB; similarly, Figure 11b shows that accuracy is stable for PD input noise up to 10^{-8} W RMS, with sharp degradation beyond, motivating noise-aware training or averaging at readout. As SNDR is a continuous metric calculated against an ideal reference using random test vectors, it degrades smoothly, whereas top-1 accuracy is a discrete metric that is resilient until the noise becomes comparable to the class-discriminative signal differences.

5.4 Waveguide Layout: Readout Region and Spatial Nonuniformity

Our waveguide configuration analysis identifies position-dependent error sources and supports component-by-component parameterization. Testing with 8- and 16-channel configurations reveals that signal location produces unique spatial error mechanisms.

We employed several spatial configurations to estimate position-dependent error. The notation is summarized in Table 2. We define these configurations using P_T, P_M, and P_B, where $P_M = \Delta - (M + N)/2$ and Δ is the center-to-center separation of the signal and kernel at the input field. The remaining channels are split such that $P_T + P_B = L - M - N - P_M$, with examples shown in Figure 12b–d. We also specify the readout region, *i.e.*, the portion of the total output field that is taken as the convolution output.

Table 4. JTC readout configurations with 8-element input signal, 8-element kernel signal, and no middle padding between input and kernel signals

Configuration	P_T	P_M	P_B	Total Output	Readout Region	Final Accuracy	Improvement
Baseline[4-12]	0	0	0	16	[4–12]	48.83%	Reference
Split[0-4,12-16]	0	0	0	16	[0–4] ∪ [12–16]	50.97%	+2.14%
Padding[8-16]	8	0	8	32	[8–16]	51.73%	+2.90%
Padding[12-20]	8	0	8	32	[12–20]	47.42%	-1.41%
Padding[16-24]	8	0	8	32	[16–24]	46.73%	-2.10%
Adjacent_Avg	8	0	8	32	Average	49.27%	+0.44%
Merged[8-24]	8	0	8	32	[8–24]	45.18%	-3.65%
Edge[0-8,24-32]	8	0	8	32	[0–8] ∪ [24–32]	43.38%	-5.45%

As an example, consider the configuration illustrated in Figure 12c. Here, both the input signal and the kernel occupy 8 channels ($M = N = 8$), with a middle gap between them ($P_M = 3$). Eight unused channels are padded above the input ($P_T = 8$) and seven below the kernel ($P_B = 7$), yielding a total lens field size of $L = P_T + M + P_M + N + P_B = 8 + 8 + 3 + 8 + 7 = 34$ channels. The readout captures the 8 output channels corresponding to the cross-correlation term, as depicted in Figure 12f.

Edge channels (0–8 and 24–32) are 7% less accurate than center channels (12–20), which points to spatially nonuniform error accumulation. Degradation arises from optical-field distribution asymmetry near the lens edge and waveguide coupling variation.

Table 4 shows that 16-channel systems have higher sensitivity to edge effects than 8-channel implementations. Central channels maintain ~50% accuracy after retraining, while edge channels drop to 43–45%. These spatial error gradients represent a fundamental error source in JTC systems, where physical position within the optical field directly correlates with achievable accuracy independent of component quality.

5.5 Correlation-Term Separation: Field Size and Signal-Kernel Spacing

Figure 12a shows the number of overlap-free cross-correlation points at the output as the lens field size and the separation between the input and kernel signals are varied. Ideally, the cross-correlation and auto-correlation terms will be completely separated in space, allowing detection without interference. However, phase noise, interference, or nonideal modulation may shift the spectra, necessitating more than the minimum separation defined in Equation 9, where M and N are the signal and kernel lengths and Δ_{min} is the minimum center-to-center spacing.

$$\Delta_{min} = M + \max(M, N) - 1 \tag{9}$$

Equation 10 describes the minimum lens size L required to avoid overlapping correlation terms.

$$L \geq 2\Delta + 2N - 1 \tag{10}$$

Equations 9 and 10 can be combined to form Equation 11, which describes the minimum lens size required based only on the input and kernel lengths.

$$L \geq 2M + 2N + 2\max(M, N) - 3 \tag{11}$$

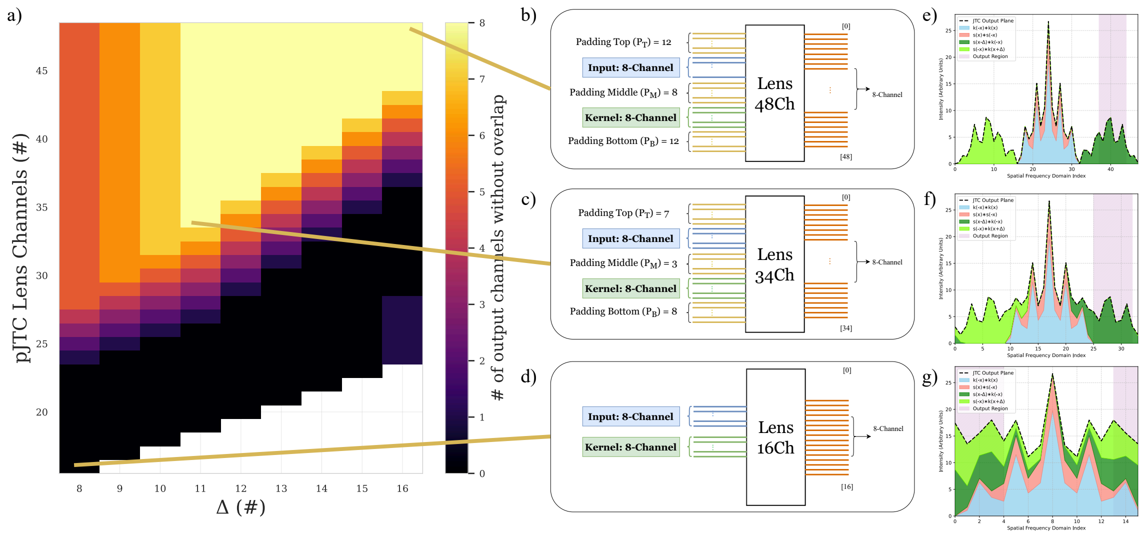

Fig. 12. a) Heatmap of the number of cross-correlation points without overlap from other correlation terms in a pJTC at different lens sizes and separations between input and kernel. b–d) Sample input/kernel configurations. e–g) Corresponding sample output correlation locations.

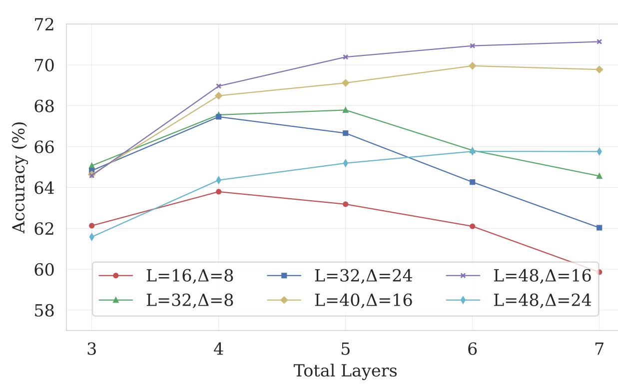

Fig. 13. Accuracy vs. network depth at different points in Figure 12a.

Figure 13 shows that separation of auto- and cross-correlation terms affects network design. The output-region overlap can be seen in Figure 12e–g. In this test case, we modulate network depth by varying the number of identical intermediate layers D before the output linear layers, and we train the network from scratch for each setup. This study only considers nonidealities from different L and Δ values to ensure that the only limitation is overlap of auto- and cross-correlation terms. When overlap is significant, ~50% of the output power of each layer comes from auto-correlation terms. As autocorrelation power dominates the output, the convolution's expressiveness decreases with each subsequent layer, and network accuracy degrades as depth increases.

6 End-to-End Accuracy Impact and Nonideality-Aware Training

Training neural networks on photonic JTC accelerators presents unique challenges stemming from hardware-enforced quantization and nonidealities distributed throughout the optical processing pipeline. Our comprehensive analysis reveals how physical limitations at each component—from electrical-to-optical conversion through Fourier-domain processing—cascade through the system, with distinct impacts on network convergence and accuracy.

Table 5. Component nonideality impact on accuracy. In this table, the baseline configuration represents an ideal photonic pipeline with transfer function $\alpha = 0$ for all blocks, no quantization, and an ideal FFT lens. "Quant only" analyzes the impact of quantization under ideal transfer functions. "All" enables quantization together with all modeled distortions. This baseline is distinct from the all-digital baselines in Table 3.

Nonideality	Inference (%)	Fine-tune (%)	Retrain (%)
Baseline	43.1	43.8	43.6
Quant only	11.4	24.2	33.4
Driver distortion	10.2	22.8	33.7
MRM amplitude distortion	10.0	24.4	32.2
MRM phase distortion	9.6	27.3	48.3
Lens distortion	10.0	32.0	46.9
PD distortion	9.8	34.5	43.9
TIA distortion	10.4	23.4	49.9
All (quant + distortions)	10.7	21.1	29.2

Table 5 summarizes the impact of hardware nonidealities on model accuracy. We evaluate the accuracy first with training on an ideal accelerator without quantization, then inference with different nonidealities enabled. We first test in a one-at-a-time configuration, then with all nonidealities applied, as would be the case in a realized accelerator. We also evaluate performance after fine-tuning for five additional epochs with a 10^{-4} learning rate, and after training from scratch with nonidealities applied throughout training. Fine-tune and retrain results are averaged over five independently seeded runs.

It can be seen that the direct application of weights trained assuming an ideal accelerator does not show good performance under severe nonidealities. For example, in the combined "All" setting, inference accuracy drops to 10.7%. Table 5 shows that fine-tuning can improve accuracy (to 21.1%), and that training from scratch with nonidealities applied throughout training can further recover accuracy (to 29.2%). To maximize performance, weights can be trained based on nonideal transfer functions (PDK-simulation-derived, or measurement-derived when available), improving accuracy relative to direct inference from Baseline-trained weights.

The retrain column in Table 5 shows that the PD, TIA, MRM phase, and lens distortions do not degrade accuracy under nonideality-aware training but rather improve accuracy relative to the baseline. These transfer functions are smooth and information-preserving, which training can leverage, and the saturating TIA response in particular acts as an additional nonlinear activation in an otherwise nearly linear optical layer. For example, when the TIA transfer function is replaced by a tanh function, the model trains to 46.6%, suggesting that the accuracy gain comes from the added smooth nonlinearity rather than the specific hardware curve. Conversely, the driver and MRM-amplitude distortions cannot be recovered by retraining. Both transfer functions flatten at the ends of their ranges, which discards input information that training is not able to restore. When the driver response is replaced with an equal-magnitude, inverted version of the driver nonlinearity, the model increases in accuracy to 47.1%, indicating that information loss rather than distortion magnitude is the primary cause of accuracy loss.

Because MRMs are thermally sensitive, we assume the system employs resonance locking and/or temperature control such that thermal drift during inference is negligible. We therefore focus on process variations and residual thermal shifts when evaluating system performance.

6.1 Fourier-Plane Sampling Quantization

The single-lens pJTC architecture requires that the JPS is squared, sampled, and digitized after the first Fourier pass, then regenerated at the input to the same lens to implement the inverse Fourier pass. A major drawback to this approach is the necessity of quantization in the Fourier plane. Figure 12e–g show that high-spatial-frequency regions can be much smaller in magnitude than low-frequency regions, even though they often contain more information. This means that if we reuse the same ADC design (and therefore bit depth and input range), we may incur larger quantization noise in the high-frequency terms of interest.

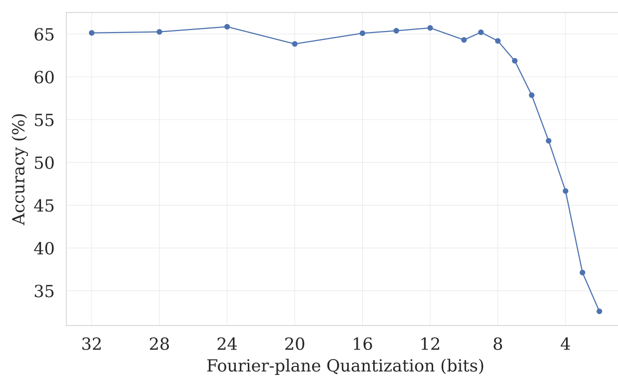

Fig. 14. Impact of Fourier-plane quantization bit depth on end-to-end model accuracy.

Figure 14 shows that increased quantization levels in the Fourier plane can have a major impact on accuracy below 8 bits of precision. Thus, when power is a primary constraint, integrating a second lens to avoid Fourier plane quantization can be more attractive than relying on a high-precision Fourier-plane digitization loop. To provide a quantitative comparison, we use the component budgets and scaling model reported in the supplementary material of [40]. Table 6 summarizes the resulting area/power breakdown for the $N_{\text{wg}} = 16$, $N_{\lambda} = 1$ case (matching the 16-channel lens used in our prototype), computed using the scaling model in [40] and the corresponding breakdown for a baud-limited single-lens case. This conservative estimate does not account for the additional Fourier plane DAC/ADC power or area required to achieve >6-bit precision, which is shown to be necessary in Figure 14.

A potential mitigation that maintains a single ADC design across all channels is to reduce the effective dynamic range seen by the ADC front end through spatial frequency-dependent optical scaling. For example, low-frequency channels where large signal magnitudes are expected could be optically attenuated to half power using waveguide splitters, reducing ADC dynamic range requirements and thus quantization noise for smaller signals. After digitization, the attenuated channel magnitudes can be bit-shifted to restore their relative values. This type of companding improves the relative precision of each channel without increasing ADC design requirements.

Table 6. First-order area/power comparison between a dual-lens pJTC (no intermediate digitization) and a single-lens digitize-and-relaunch variant for $N_{\text{wg}} = 16$, $N_\lambda = 1$.

	Dual-lens	**Single-lens + digitize/relaunch**
Power (W)		
ADC(s)	0.309	0.309
DAC(s)	0.0711	0.0711
MRR(s)	0.0320	0.0160
Laser(s)	0.0120	0.0120
Total	**0.424**	**0.408**
Area (mm^2)		
Lens(es)	7.68×10^{-2}	3.84×10^{-2}
ADC(s)	3.04×10^{-2}	3.04×10^{-2}
DAC(s)	2.44×10^{-2}	2.44×10^{-2}
PD(s)	3.87×10^{-2}	1.93×10^{-2}
MRR(s)	1.57×10^{-2}	7.84×10^{-3}
Beam splitters	5.41×10^{-3}	5.41×10^{-3}
Total	$\mathbf{1.91 \times 10^{-1}}$	$\mathbf{1.26 \times 10^{-1}}$
Derived metrics (baud-limited digitize-and-relaunch)		
Throughput (TOPS)	1.28	0.640
TOPS/W	3.02	1.57
TOPS/mm^2	6.69	5.09

7 Conclusion

We present a flexible digital twin of a SiPho JTC system to quantify how hardware nonidealities impact convolution and CNN accuracy. This enables sensitivity-driven design-space exploration to identify the most critical hardware improvements without repeatedly running device-level simulations. Our analysis shows that lens distortion is the most damaging distortion at inference time without retraining as it changes the Fourier mixing behavior rather than acting as a distribution shift that can be corrected by batch-normalization statistics recalibration or nonideality-aware training. With retraining, accuracy is limited by MRM amplitude nonlinearity and driver distortion, whereas TIA saturation, parasitic phase, and lens distortion train to accuracies at or above the ideal-hardware baseline due to increased nonlinearity-induced expressivity. We also observe that capturing a non-overlapped cross-correlation term is key to network performance: both selecting the correct output channels and ensuring that they are not overlapped affect accuracy, especially as network depth increases. These findings help guide future designs, where we size the lens field and choose input/kernel separation to satisfy the derived minimum conditions, and we prioritize minimizing lens, driver, and MRM amplitude distortions.

In general, the framework helps identify which areas in the system benefit most from hardware performance improvements. After design improvements, updated transfer functions can be swapped into the digital twin to validate end-to-end accuracy impact and iterate toward performance targets. Our current study uses steady-state, memoryless block transfer functions, augmented with effective stochastic laser/PD noise models and a reduced-order unitary lens distortion model. Extending the framework to capture dynamic bandwidth/transient effects, as well as additional noise sources, is left for future work.

References

[1] M Alam and M Karim. 1993. Fringe-adjusted joint transform correlation. *Appl. Opt.* 32, 23 (Aug. 1993), 4344–4350. doi:10.1364/AO.32.004344

[2] Sanmitra Banerjee, Mahdi Nikdast, and Krishnendu Chakrabarty. 2021. Modeling silicon-photonic neural networks under uncertainties. In *2021 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 98–101. doi:10.23919/DATE51398.2021.9474000

[3] Sanmitra Banerjee, Mahdi Nikdast, and Krishnendu Chakrabarty. 2023. On the Impact of Uncertainties in Silicon-Photonic Neural Networks. *IEEE Design & Test* 40, 2 (2023), 82–89. doi:10.1109/MDAT.2022.3161599

[4] Santiago Bernal, Mario Dumont, Essam Berikaa, Charles St-Arnault, Yixiang Hu, Ramon Gutierrez Castrejon, Weijia Li, Zixian Wei, Benjamin Krueger, Fabio Pittalà, et al. 2024. 12.1 terabit/second data center interconnects using O-band coherent transmission with QD-MLL frequency combs. *Nature Communications* 15, 1 (2024), 7741. doi:10.1038/s41467-024-51682-2

[5] Yusheng Bian, Takako Hirokawa, Won Suk Lee, Sujith Chandran, Ken Giewont, Abdelsalam Aboketaf, Qidi Liu, Ryan Sporer, Michal Rakowski, Kevin Dezfulian, et al. 2024. 300-mm monolithic CMOS silicon photonics foundry technology. In *CLEO 2024 (Technical Digest Series)*. Optica Publishing Group, ATu3H.1. doi:10.1364/CLEO_AT.2024.ATu3H.1 Invited paper.

[6] Chia-Ming Chang, Po Dong, Chengcheng Gui, and Guilhem de Valicourt. 2016. Low-chirp Push-pull Microring Modulators. In *Optical Fiber Communication Conference*. Optica Publishing Group, Th4H.2. doi:10.1364/OFC.2016.Th4H.2

[7] Xue-Wen Chen. 1998. Distortion-invariant fractional power fringe-adjusted joint transform correlation. *Opt. Eng.* 37, 1 (Jan. 1998), 138. doi:10.1117/1.601840

[8] Raphaël Dubé-Demers, Sophie LaRochelle, and Wei Shi. 2016. Ultrafast pulse-amplitude modulation with a femtojoule silicon photonic modulator. *Optica* 3, 6 (2016), 622–627. doi:10.1364/OPTICA.3.000622

[9] Sally I El-Henawy, Zhengxing Zhang, Ryan Miller, and Duane S Boning. 2019. Photonic device sensitivity analysis methods: towards process variation-aware silicon photonics design. In *Optical Modeling and System Alignment*, Vol. 11103. SPIE, 185–195. doi:10.1117/12.2529198

[10] Alireza Geravand, Zibo Zheng, Simon Levasseur, Leslie A Rusch, and Wei Shi. 2023. Ultra-compact silicon modulator for 124 GBaud coherent optical links. In *2023 IEEE Silicon Photonics Conference (SiPhotonics)*. IEEE, 1–2. doi:10.1109/SiPhotonics55903.2023.10141908

[11] Alireza Geravand, Zibo Zheng, Farshid Shateri, Simon Levasseur, Leslie A Rusch, and Wei Shi. 2025. Ultrafast coherent dynamics of microring modulators. *Nature Photonics* 19, 7 (2025), 740–750. doi:10.1038/s41566-025-01686-1

[12] Joseph W Goodman and Mary E Cox. 1969. *Introduction to Fourier optics.* American Institute of Physics.

[13] Thomas J Grycewicz. 1999. Techniques to improve binary joint transform correlator performance for fingerprint recognition. *Opt. Eng.* 38, 1 (Jan. 1999), 114. doi:10.1117/1.602066

[14] Nicholas J Higham. 2008. *Functions of Matrices: Theory and Computation.* SIAM Publications, Philadelphia, PA. doi:10.1137/1.9780898717778

[15] H Inbar, N Konforti, and E Marom. 1994. Modified joint transform correlator binarized by error-diffusion. I. Spatially constant noise-dependent range limit. *Appl. Opt.* 33, 20 (July 1994), 4434–4443. doi:10.1364/AO.33.004434

[16] H Inbar, D Mendlovic, and E Marom. 1993. Error-diffusion binarization for joint transform correlators. *Appl. Opt.* 32, 5 (Feb. 1993), 707–714. doi:10.1364/AO.32.000707

[17] E. Kewei, Chen Zhang, Mengyang Li, Zhao Xiong, and Dahai Li. 2015. Wavefront reconstruction algorithm based on Legendre polynomials for radial shearing interferometry over a square area and error analysis. *Opt. Express* 23, 16 (Aug. 2015), 20267–20279. doi:10.1364/OE.23.020267

[18] Jiachen Li, Sigang Yang, Hongwei Chen, Xingjun Wang, Minghua Chen, and Weiwen Zou. 2022. Fully integrated hybrid microwave photonic receiver. *Photonics Research* 10, 6 (2022), 1472–1483. doi:10.1364/PRJ.452631

[19] Rui Li, D Patel, E El-Fiky, Alireza Samani, Zhenping Xing, M Morsy-Osman, and D Plant. 2017. High-speed low-chirp PAM-4 transmission based on push-pull silicon photonic microring modulators. *Opt. Express* 25, 12 (June 2017), 13222–13229. doi:10.1364/OE.25.013222

[20] Shurui Li, Hangbo Yang, Chee Wei Wong, Volker J Sorger, and Puneet Gupta. 2023. Photofourier: A photonic joint transform correlator-based neural network accelerator. In *2023 IEEE International Symposium on High-Performance Computer Architecture (HPCA)*. IEEE, 15–28. doi:10.1109/HPCA56546.2023.10070931

[21] Yibo Li, Jiang Guo, and Rengcong Liu. 2024. Exploring Wavefront Detection in Imaging Systems with Rectangular Apertures Using Phase Diversity. *Sensors* 24, 4 (2024), 1191. doi:10.3390/s24041191

[22] Yinyi Liu, Jiaxu Zhang, Jun Feng, Shixi Chen, and Jiang Xu. 2022. A reliability concern on photonic neural networks. In *2022 Design, Automation & Test in Europe Conference & Exhibition (DATE)*. IEEE, 1059–1064. doi:10.23919/DATE54114.2022.9774661

[23] Yanir London, Thomas Van Vaerenbergh, Anthony J Rizzo, Peng Sun, Jared Hulme, Geza Kurczveil, Ashkan Seyedi, Binhao Wang, Xiaoge Zeng, Zhihong Huang, et al. 2020. Energy efficiency analysis of comb source carrier-injection ring-based silicon photonic link. *IEEE Journal of Selected Topics in Quantum Electronics* 26, 2 (2020), 1–13. doi:10.1109/JSTQE.2019.2934121

[24] G Lu, Z Zhang, S Wu, and F T Yu. 1997. Implementation of a non-zero-order joint-transform correlator by use of phase-shifting techniques. *Appl. Opt.* 36, 2 (Jan. 1997), 470–483. doi:10.1364/AO.36.000470

[25] AV Masalov, A Kuzhamuratov, and AI Lvovsky. 2017. Noise spectra in balanced optical detectors based on transimpedance amplifiers. *Review of Scientific Instruments* 88, 11 (2017). doi:10.1063/1.5004561

[26] Alvaro Moscoso-Mártir, Juliana Müller, Johannes Hauck, Nicolas Chimot, Rony Setter, Avner Badihi, Daniel E. Rasmussen, Alexandre Garreau, Mads Nielsen, Elmira Islamova, Sebastián Romero-García, Bin Shen, Anna Sandomirsky, Sylvie Rockman, Chao Li, Saeed Sharif Azadeh, Guo-Qiang Lo, Elad Mentovich, Florian Merget, François Lelarge, and Jeremy Witzens. 2017. Silicon Photonics Transmitter with SOA and Semiconductor Mode-Locked Laser. *Scientific Reports* 7, 1 (2017), 13857. doi:10.1038/s41598-017-14347-3

[27] Riddhi Nandi, Pratyasha Priyadarshini, Crystal Hedges, Frank Pavlik, and Michal Rakowski. 2024. Photonic integrated circuit design flow in 45 nm SOI GF Fotonix™ platform. In *NUSOD 2024*. https://www.nusod.net/wp-content/uploads/2024/abstracts/NUSOD24_WB01.pdf

[28] Shupeng Ning, Hanqing Zhu, Chenghao Feng, Jiaqi Gu, Zhixing Jiang, Zhoufeng Ying, Jason Midkiff, Sourabh Jain, May H Hlaing, David Z Pan, et al. 2024. Photonic-electronic integrated circuits for high-performance computing and ai accelerators. *Journal of Lightwave Technology* 42, 22

(2024), 7834–7859. doi:10.1109/JLT.2024.3427716

[29] Quantifi Photonics Ltd. 2022. Relative Intensity Noise (RIN). Application note, Version 1.1. https://cdn.quantifiphotonics.com/20220527144951/QP_AppNote_RIN_v1.1.pdf

[30] Michal Rakowski, Colleen Meagher, Karen Nummy, Abdelsalam Aboketaf, Javier Ayala, Yusheng Bian, Brendan Harris, Kate Mclean, Kevin McStay, Asli Sahin, Louis Medina, Bo Peng, Zoey Sowinski, Andy Stricker, Thomas Houghton, Crystal Hedges, Ken Giewont, Ajey Jacob, Ted Letavic, Dave Riggs, Anthony Yu, and John Pellerin. 2020. 45nm CMOS - Silicon Photonics Monolithic Technology (45CLO) for next-generation, low power and high speed optical interconnects. In *Optical Fiber Communication Conference (OFC) 2020 (OSA Technical Digest)*. Optica Publishing Group, T3H.3. doi:10.1364/OFC.2020.T3H.3

[31] Eduard Sackinger. 2005. *Broadband circuits for optical fiber communication*. John Wiley & Sons, Nashville, TN. doi:10.1002/0471726400

[32] Charbel Sakr, Steve Dai, Rangha Venkatesan, Brian Zimmer, William Dally, and Brucek Khailany. 2022. Optimal clipping and magnitude-aware differentiation for improved quantization-aware training. In *International Conference on Machine Learning*. PMLR, 19123–19138. https://proceedings.mlr.press/v162/sakr22a.html

[33] Amin Shafiee, Sanmitra Banerjee, Krishnendu Chakrabarty, Sudeep Pasricha, and Mahdi Nikdast. 2024. Analysis of Optical Loss and Crosstalk Noise in MZI-Based Coherent Photonic Neural Networks. *Journal of Lightwave Technology* 42, 13 (2024), 4598–4613. doi:10.1109/JLT.2024.3373250

[34] H J Su and M A Karim. 1999. Phase-shifting joint transform correlation with phase-iterative algorithm: effect of the dynamic range limit. *Appl. Opt.* 38, 26 (Sept. 1999), 5556–5559. doi:10.1364/AO.38.005556

[35] Qing Tang and Bahram Javidi. 1992. Sensitivity of the nonlinear joint transform correlator: experimental investigations. *Appl. Opt.* 31, 20 (Jul 1992), 4016–4024. doi:10.1364/AO.31.004016

[36] Isidora Teofilovic, Ali Cem, David Sanchez-Jacome, Daniel Pérez-López, and Francesco Da Ros. 2024. Thermal crosstalk modelling and compensation methods for programmable photonic integrated circuits. *Journal of Lightwave Technology* 42, 22 (2024), 7816–7824. doi:10.1109/JLT.2024.3430504

[37] Vincent J Urick, John F Diehl, Meredith N Draa, Jason D McKinney, and Keith J Williams. 2012. Wideband analog photonic links: some performance limits and considerations for multi-octave implementations. In *Rf and Millimeter-Wave Photonics Ii*, Vol. 8259. Spie, 26–39. doi:10.1117/12.912528

[38] R Wang, L Shang, and C Chatwin. 1996. Modified fringe-adjusted joint transform correlation to accommodate noise in the input scene. *Appl. Opt.* 35, 2 (Jan. 1996), 286–296. doi:10.1364/AO.35.000286

[39] X Wu, B Guan, Qianfan Xu, C Doerr, and Long Chen. 2020. Low-chirp push-pull dual-ring modulator with 144 Gb/s PAM-4 data transmission. *Opt. Express* 28, 18 (Aug. 2020), 26492–26498. doi:10.1364/OE.399800

[40] Hangbo Yang, Nicola Peserico, Shurui Li, Xiaoxuan Ma, Russell L. T. Schwartz, Mostafa Hosseini, Aydin Babakhani, Chee Wei Wong, Puneet Gupta, and Volker J. Sorger. 2025. Near-energy-free photonic Fourier transformation for convolution operation acceleration. *Advanced Photonics* 7, 5 (2025), 056007. doi:10.1117/1.AP.7.5.056007

[41] Lin Zhang, Jeng-Yuan Yang, Muping Song, Yunchu Li, Bo Zhang, Raymond G Beausoleil, and Alan E Willner. 2007. Microring-based modulation and demodulation of DPSK signal. *Optics express* 15, 18 (2007), 11564–11569. doi:10.1364/OE.15.011564

[42] Zhengxing Zhang, Milica Notaros, Zhengqi Gao, Uttara Chakraborty, Jelena Notaros, and Duane S Boning. 2023. Impact of process variations on splitter-tree-based integrated optical phased arrays. *Opt. Express* 31, 8 (April 2023), 12912–12921. doi:10.1364/OE.487096

[43] Dongsheng Zheng, Ciyuan Qiu, Hongxia Zhang, Xinhong Jiang, and Yikai Su. 2020. Demonstration of a push-pull silicon dual-ring modulator with enhanced optical modulation amplitude. *J. Lightwave Technol.* 38, 14 (July 2020), 3694–3700. doi:10.1109/JLT.2020.2978149

[44] Zibo Zheng, Xun Guan, Alireza Geravand, Wei Shi, and Leslie A Rusch. 2024. Enhanced performance for SSB-OFDM transmission with a low-chirp silicon MRM. *J. Lightwave Technol.* 42, 9 (2024), 1–10. doi:10.1109/JLT.2024.3362248