Improving the Reliability and Energy-Efficiency of High-Bandwidth Photonic NoC Architectures with Multilevel Signaling

Ishan G Thakkar, Sai Vineel Reddy Chittamuru, Sudeep Pasricha
Department of Electrical and Computer Engineering
Colorado State University, Fort Collins, CO, U.S.A.
{ishan.thakkar, sai.chittamuru, sudeep}@colostate.edu

ABSTRACT

Photonic network-on-chip (PNoC) architectures employ photonic waveguides with dense-wavelength-division-multiplexing (DWDM) for signal traversal and microring resonators (MRs) for on-off-keying (OOK) based signal modulation, to enable high bandwidth on-chip transfers. Unfortunately, the use of larger number of DWDM wavelengths to achieve higher bandwidth requires sophisticated and costly laser sources along with extra photonic hardware, which adds extra noise and increases the power and area consumption of PNoCs. This paper presents a novel method (called *4-PAM-P*) of generating four-amplitude-level optical signals in PNoCs, which doubles the aggregate bandwidth without increasing utilized wavelengths, photonic hardware, and incurred noise, thereby reducing the bit-error-rate (BER), area, and energy consumption of PNoCs. Our experimental analysis shows that our *4-PAM-P* signaling method achieves equal bandwidth with 4.2× better BER, 19.5% lower power, 16.3% lower energy-per-bit, and 5.6% less photonic area compared to the best known 4-amplitude-level optical signaling method from prior work.

Categories and Subject Descriptors: [Networks] Network on chip; [Hardware] Integrated Circuits/Interconnect: Photonic and optical interconnect

Keywords: Photonic network on chip; multilevel optical signaling; optimization; energy efficiency; reliability

1 INTRODUCTION

In the many-core era, processors with hundreds of cores on a single chip are gradually becoming a reality. The performance of these many-core processors is driven by the effectiveness of the underlying electrical network-on-chip (ENoC) fabrics that are becoming increasingly crosstalk- and energy-limited [1]. To this end, due to the recent developments in the area of silicon photonics, photonic network-on-chip (PNoC) fabrics have been projected to supersede ENoCs. PNoCs offer several benefits over ENoCs such as higher bandwidth density, distance-independent bit-rate, and smaller data-dependent energy.

NOCS '17, October 19–20, 2017, Seoul, Republic of Korea

ACM ISBN 978-1-4503-4984-0/17/10…$15.00
https://doi.org/10.1145/3130218.3130226

Typical PNoC architectures (e.g., [1]-[4], [8]) and off-chip photonic interconnects (e.g., [25], [26]) utilize several photonic devices such as multi-wavelength lasers, microring resonators (MRs), waveguides (WGs), and splitters. A broadband laser source generates light of multiple wavelengths (λs), each wavelength (λ) of which serves as a data signal carrier. Simultaneous traversal of multiple λ-signals across a single photonic WG is possible using dense wavelength division multiplexing (DWDM), which enables parallel data transfer across the photonic WG. For instance, a DWDM of 64 λs can transfer 64 data bits in parallel. At the source node, multiple MRs typically modulate multiple electrical data signals on the utilized DWDM λs (data-modulation phase). In almost all PNoCs in literature, modulator MRs utilize on-off keying (OOK) modulation, where-in the presence and absence of λs in the WG are used to represent logic '1s' and '0s'. Similarly, at the destination node, multiple MRs with photodetectors at their drop ports are used to filter and detect light-modulated data signals from the WG (data-detection phase) and generate proportional electrical signals. In general, the use of a large number of DWDM λs enables high bandwidth parallel data transfers in PNoCs.

Unfortunately, a number of challenges related to area [8], cost [32], reliability [11], and energy-efficiency [13] still need to be overcome for efficient implementation of PNoCs that utilize a large number of DWDM λs (typically 64 or more DWDM λs per WG). Generating a large number of DWDM λs requires a comb-generating laser source, the ineffectiveness, complexity, and cost of which increase with the number of λs generated [5]. Moreover, utilizing a larger number of DWDM λs to achieve higher-bandwidth data-transfers in a PNoC results in larger network flit size and more electrical and photonic hardware (more number of modulator and detector MRs and their drivers). A larger network flit size also results in larger sized buffers in the network gateway interfaces, which results in significantly higher area and power overheads. Similarly, larger number of MRs and drivers also incur greater photonic area and MR heating power overheads. Furthermore, the use of a larger number of DWDM λs decreases the gap between two successive λ-channels, which in turn increases the heterodyne crosstalk noise in PNoCs, harming the reliability of communication [6], [9]. Thus, the use of larger number of DWDM λs to achieve higher bandwidth in PNoCs is not a reliable and energy-efficient option. *This motivates the need for a more reliable and energy-efficient way of achieving higher bandwidth data transfers in PNoC architectures*.

In [4] and [7], Kao et al. proposed a multilevel optical signaling format 4-PAM (4-pulse amplitude modulation) to achieve higher bandwidth and energy-efficient data communication in PNoCs. 4-PAM optical signaling format doubles the bandwidth by compressing two bits in one symbol carried out by four levels of amplitude. Kao et al. utilize superposition of two OOK-modulated optical signals of the same λ with 2:1 power ratio to create a 4-PAM λ-signal. We found that this signal superposition based 4-PAM optical signaling method (referred to as 4-PAM-SS henceforth)

incurs significantly high power, photonic area, and reliability overheads (Section II). *The shortcomings of 4-PAM-SS motivate the need for a more reliable, energy- and area-efficient method of implementing 4-PAM signaling*.

In this paper, we present a novel method (referred to as *4-PAM-P* henceforth) of generating 4-PAM optical signals, which employs only one modulator MR per λ to directly modulate the designated λ-signal in 4-PAM format. We present a search heuristic based link optimization framework that finds the optimal value of number of DWDM wavelengths (N_λ) from a constrained space of all its allowable values to achieve desired performance and/or reliability goals for the target photonic link. We use our framework to optimize the designs of three types of photonic links, each of which uses OOK, 4-PAM-SS, or *4-PAM-P* optical signaling method. Our experimental analysis shows that PNoCs that are comprised of *4-PAM-P* based optimized links render greater reliability, energy-efficiency, and area-efficiency with equal bandwidth compared to the PNoCs that are comprised of 4-PAM-SS or OOK based optimized links. We summarize the key contributions in this paper as follows:

- We propose a novel technique (*4-PAM-P*) of generating 4-PAM optical signals in PNoCs, which proves to be more reliable, area- and energy-efficient than the previously proposed signal superposition based 4-PAM-SS method [7] and the conventional OOK method;
- We present a search heuristic based optimization framework that optimizes the designs of OOK, 4-PAM-SS, and *4-PAM-P* signaling based photonic links to achieve desired performance and/or reliability goals;
- We evaluate the impact of the optimized designs of OOK, 4-PAM-SS and *4-PAM-P* based photonic links on the performance, reliability and energy-efficiency of a well-known PNoC architecture: an 8-ary 3-stage CLOS PNoC [4].

2 BACKGROUND AND MOTIVATION

In this section, we first present an overview of the signal superposition based 4-PAM optical signaling method (referred to as 4-PAM-SS) from [7]. In 4-PAM-SS, at first, two separate OOK-modulated signals of each of the utilized DWDM λs are generated in two separate parallel WGs. Then, these two sets of OOK-modulated DWDM signals are superposed using a combiner to generate one set of 4-PAM modulated DWDM signals. Each amplitude level of a 4-PAM λ-signal represents one of the four combinations of two bits (00, 01, 10, or 11). When the entire DWDM spectrum of 4-PAM signals reach the destination node, each signal is filtered by its corresponding in-resonance MR and is converted back into two electrical signals by a photodetector and a back-end receiver circuit. As discussed in [7], the back-end receiver circuit consists of three sense-amplifiers and two logic gates that decode the 4-PAM modulated signal.

Ideally, in the 4-PAM-SS method, when the combiner superposes two OOK-modulated signals, a 4-PAM modulated signal is generated owing to the constructive interference between the two OOK signals. However, the constructive interference happens only if both the OOK signals have identical phases. Unfortunately, in the presence of non-idealities such as process and on-chip temperature variations, a significant phase difference exists between the two superposed OOK signals, which leads to destructive interference between them. Owing to the random nature of process and on-chip temperature variations, this incurred phase difference may fall anywhere in the range from 0 to 2π. This implies that the degree of destructive interference incurred between the OOK signals due to the phase difference (and hence the strengths of the symbols of the resultant 4-PAM signal) may fall anywhere in a very large ranges of values. This in turn makes it very hard to ensure reliability of communication with a 4-PAM-SS photonic link.

The worst-case destructive interference in 4-PAM-SS occurs when the two superposed OOK signals are completely out of phase, i.e., when the phase difference between them is an odd multiple of π. The amount of signal loss due to the superposition of two out of phase OOK signals depends on their individual signal strengths. Typically, as explained in [7], in 4-PAM-SS method, to equidistantly space the four amplitude levels of the output 4-PAM signal in the available range of optical transmission, the strengths of the individual OOK signals are kept to be two-third and one-third of the strength of the conventional OOK signal. Hence, for the best-case constructive interference between the superposed OOK signals, the strength of the resultant 4-PAM signal becomes 2/3+1/3=1. In contrast, for the worst-case destructive interference, the strength of the resultant 4-PAM signal becomes 2/3-1/3=1/3, which makes the worst-case interference-related signal loss to be -10×log(1/3) = 4.8dB.

In summary, the interference-related signal loss in 4-PAM-SS [7] reduces signal-to-noise ratio (SNR), BER, and overall communication reliability. Furthermore, as explained in [7], the 4-PAM-SS method requires additional photonic hardware, such as one asymmetric splitter, two modulator MRs per λ (for two OOK signals), and a combiner. This additional photonic hardware reduces the area benefits of the 4-PAM-SS method compared to the traditional OOK method. In this paper, we present a novel, more reliable, and energy- and area-efficient 4-PAM signaling method (*4-PAM-P*), which overcomes the shortcomings of the 4-PAM-SS method. The next section describes our proposed *4-PAM-P* method in detail.

3 PROPOSED 4-PAM-P OPTICAL SIGNALING

3.1 Overview

Unlike 4-PAM-SS, our proposed *4-PAM-P* method employs only one modulator MR per λ to directly modulate the designated λ-signal in 4-PAM format. This type of 4-PAM signal generating modulator MRs are demonstrated in [19], [21], and [22]. Our proposed *4-PAM-P* method extends the use of such modulator MRs in DWDM based PNoCs. Before we discuss about how our *4-PAM-P* method works, it is important to understand how a modulator MR works in the conventional OOK method.

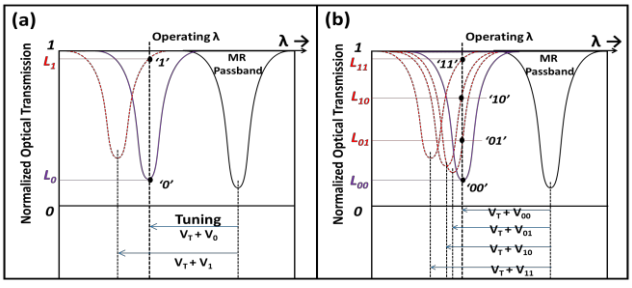

Fig. 1. Illustration of optical transmission and microring resonator (MR) spectra for (a) OOK signaling, (b) proposed *4-PAM-P* signaling.

Ideally, a modulator MR is designed to operate in resonance with a signal-λ in its default state. But due to process and on-chip temperature variations, the MR's resonance λ often deviates from the signal-λ (black curve in Fig. 1(a)). In this case, as shown in Fig. 1(a), the MR's resonance λ (center/peak of the MR's passband) needs to be brought in alignment with the signal-λ by either

temperature or electrical tuning of the MR [29]. In this tuned state (purple curve), the MR remains in resonance with the signal-λ, which enables the MR to modulate logic 0 on the signal-λ, by removing the signal-λ from the WG. Thus, in this tuned state of the MR, a non-zero tuning-bias voltage (V_T) but zero signal-bias voltage ($V_0 = 0$) is applied to the MR. Hence, as shown in Fig. 1(a), the total bias voltage applied to the MR in the tuned state (purple curve) is $V_B = V_T + V_0 \rightarrow V_B = V_T$. On the other hand, the MR is operated in the off-resonance state to modulate logic 1 on the signal-λ. For that, a specific non-zero signal-bias voltage V_1 is applied to the MR, which shifts the passband of the MR to be in off-resonance state (red curve). Thus, the net bias voltage applied in the off-resonance state of the MR is $V_B=V_T+V_1$.

In summary, in conventional OOK modulation, to modulate a particular sequence of 1s and 0s on the signal-λ, the modulator MR is switched off and on resonance with the signal-λ by applying the net bias voltages of V_T+V_1 and V_T+V_0 to the MR, respectively. Note that the value and polarity of V_T depends on the amount and direction of the variation-induced resonance shift, which can be efficiently assessed by using the dithering signal based method demonstrated in [18]. Moreover, in the OOK method (Fig. 1(a)), two levels of optical transmissions (L_0 and L_1) are achieved that correspond to net bias voltages of V_T+V_0 and V_T+V_1. The difference between L_0 and L_1 is defined as modulation depth.

In our *4-PAM-P* method, as shown in Fig. 1(b), we introduce two more intermediate levels of optical transmissions L_{01} and L_{10} between L_0 and L_1 within the modulation depth. For that, we introduce two more levels of signal-bias voltages V_{01} and V_{10} between V_0=0 and V_1. Thus, in *4-PAM-P*, signal-bias voltages/optical transmission levels V_0/L_0 (or V_{00}/L_{00}), V_1/L_1 (or V_{11}/L_{11}), V_{01}/L_{01}, and V_{10}/L_{10} correspond to bit combinations "00", "11", "01", and "10" respectively. Note that as demonstrated in [19], the resonance passbands of the carrier-depletion based optimized MRs can be shifted with a signal voltage efficiency of 2GHz/V. This allows MRs with bandwidths even as low as ~10GHz (corresponding to the quality factor of 18000) to have very fine and efficient control of optical transmission levels in the resultant 4-PAM signals [19].

In contrast to the 4-PAM-SS method, *4-PAM-P* does not use signal superposition to create 4-PAM signals, and therefore does not incur interference-related signal loss. As a result, for given noise power, *4-PAM-P* renders greater SNR with better BER. Moreover, *4-PAM-P* requires one less modulator MR per λ, and it does not require additional splitters and combiners. As a result, *4-PAM-P* consumes less photonic area and dissipates less static power related to tuning of MRs. Due to all these benefits, *4-PAM-P* is more reliable, energy- and area-efficient than the 4-PAM-SS method.

However, compared to the OOK and 4-PAM-SS methods, the use of *4-PAM-P* requires some minor modifications in how electrical-to-optical (E/O) conversion of data at the sender/modulator side is implemented. At the receiver side, optical-to-electrical (O/E) conversion of data in *4-PAM-P* is carried out in the same manner as in the 4-PAM-SS method, as demonstrated in [7]. Note that from [20] and [7], a 4-PAM signal with the same baud-rate as of an OOK signal requires 4.8dB more received power to achieve the same BER as achievable by the OOK signal. Therefore, for our link- and architecture-level analysis in Section IV and V, we consider the detector sensitivity threshold for both the 4-PAM-SS and *4-PAM-P* methods to be 4.8dB more than the conventional OOK method, as manifested by the SNR_{Target} value in Eq. (3). The next subsection describes how E/O conversion is implemented in our *4-PAM-P* method.

3.2 E/O Conversion in 4-PAM-P Method

In a typical PNoC, at the electrical-optical interface of a sender node, the input electrical flits are temporarily stored in first-in-first-out (FIFO) buffers before the modulator MRs convert them into the optical domain. Typically, for PNoCs with equal-sized electrical and optical flits, the size of each entry of the FIFO buffer is equal to the size of the optical flit, which is equal to the number of bits transferred in parallel on DWDM λs [3]. N parallel electrical bit-streams can be produced when multiple N-bit entries of the FIFO buffer are evicted in sequence, triggered by consecutive clock edges (or levels). The modulator MRs at the E/O interface convert these parallel electrical bit-streams into parallel optical bit-streams.

Fig. 2(a) illustrates E/O conversion for the conventional OOK method, using an example photonic link with 4-bit flit-size. Each of the four parallel electrical bit-streams available is mapped to a designated modulator MR that is designed to operate on a particular λ. Accordingly, each of these bit-streams is applied to a driver circuit, which produces a corresponding sequence/stream of signal-bias voltages V_0 and V_1. Then, the designated modulator MR is driven by this sequence of signal-bias voltages, after the voltage sequence is offset with a constant tuning-bias voltage V_T corresponding to and adjusting for the variation-induced resonance shift in the modulator MR. As explained in the last subsection, because of this applied sequence of bias voltages, each modulator MR (corresponding to an electrical bit-stream) modulates the input electrical bit-stream onto the corresponding DWDM λ, generating an OOK-modulated optical signal.

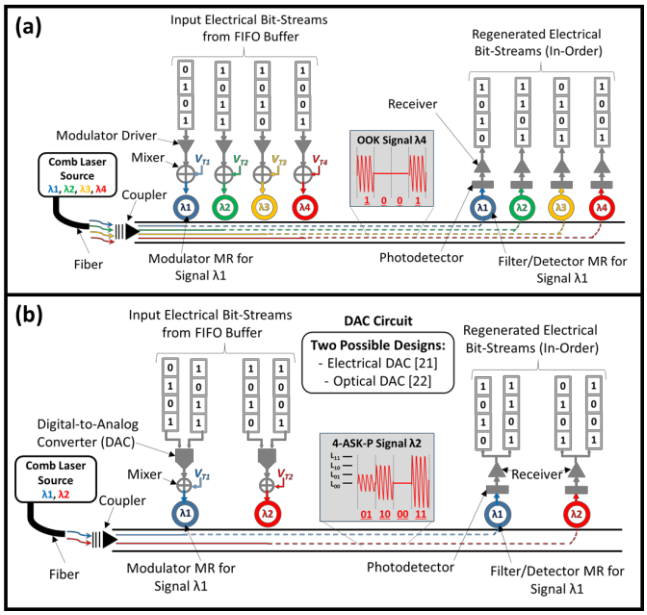

Fig. 2. Schematics of (a) OOK-based, (b) *4-PAM-P* based photonic links.

On the other hand, as shown in Fig. 2(b), in the proposed *4-PAM-P* method as well, the FIFO buffer generates four parallel electrical bit-streams. Two adjacent electrical bit-streams of these four bit-streams are applied as inputs to a digital-to-analog converter (DAC) circuit. The DAC coverts each input two-bit combination to a signal-bias voltage level out of four possible voltage levels V_{00}, V_{01}, V_{10}, and V_{11} (see Section III.A). Thus, four parallel electrical bit-streams are converted in two parallel sequences of signal-bias voltages by two concurrently operated DAC units. These two parallel sequences of signal-bias voltages, after being offset by

corresponding tuning-bias voltages, are applied to two designated modulator MRs that are designed to operate on two different λs. These modulator MRs modulate the applied sequences of four-level voltages onto their corresponding λs to generate two parallel four-amplitude-level (4-PAM) optical signals.

Now, as given in the inset of Fig. 2(b), two different designs of a high-speed DAC circuit are possible: electrical DAC [21] and optical DAC [22]. In [21], a 40Gbps DAC designed in 65nm CMOS is demonstrated, which utilizes a segmented pulsed-cascode output stage to achieve 4-PAM modulation on a single MR. This electrical DAC (demonstrated in [21]) consumes 3.04pJ/bit power to convert two input electrical bit-streams into a four-level electrical signal (4-PAM electrical signal) at 40Gbps. This 4-PAM electrical signal, after being properly conditioned by signal-bias voltages, is applied to an MR that generates a proportional optical 4-PAM signal. Thus, the conversion of two input electrical bit-streams into a single 4-PAM optical signal happens in two stages. In the first stage, the input bit-streams are converted into an electrical 4-PAM signal by the DAC circuit. Then in the second stage, this 4-PAM electrical signal is converted into optical domain by the modulator MR. The caveat of this electrical DAC based conversion method is that it incurs significant area and energy overhead and imposes non-linearity onto the E/O transfer functions of the driven MRs [22]. To overcome these shortcomings, Moazeni et al. in [22] utilized an optical DAC, which is basically a "spoked" MR of 5μm radius with ~10000 Q that directly converts two input electrical bit-streams into a 4-PAM optical signal for only 0.197pJ/bit power consumption at 20Gbps bit-rate. Thus, the use of these "spoked" MRs collapses the two-stage E/O conversion process into a single stage process, and thus, eliminates the need for external electrical DAC circuit. In this case, the DAC circuits shown in Fig. 2(b) are eliminated and the input electrical bit-streams are directly applied to the corresponding "spoked" MRs. We utilize these optical DACs ("spoked" MRs) in our 4-PAM-P method and account for 0.197pJ/bit power consumption (as part of Tx/Rx power) per DAC in our architecture-level evaluation presented in Section V.

4 PHOTONIC-LINK DESIGN METHODOLOGY

A naive design of photonic links can result in suboptimal values of bandwidth, power, and reliability for the associated PNoC [10]. Therefore, irrespective of the utilized optical signaling method, it becomes imperative to optimize the designs of the constituent photonic links to achieve maximum bandwidth, energy-efficiency, and reliability at the PNoC architecture level. From [6] and [10], for photonic-link design optimization, the number of DWDM λs per waveguide (N_λ) is the most important design parameter, and link power-budget ($\mathbf{P^{dB}_{Budget}}$) is the most critical design constraint. For optimal use of the available power budget and link bandwidth [10], parameters N_λ and $\mathbf{P^{dB}_{Budget}}$ should meet conditions given in Eq. (1), where the expressions for the constituent terms of Eq. (1) are given in Eq. (2)-(4).

Eq. (3) is derived from the equation for the required photodiode power given in [12]. In Eq. (3), *S* represents the required detector sensitivity threshold (i.e., minimum detectable power) in dBm to achieve the target SNR (SNR_{Target}). Table 1 gives the definitions and values of various parameters used in Eq. (2)-(4). As evident from Table 1, SNR_{Target} is different for OOK and 4-PAM signals, the reason for which is explained in Section IV-B. We use the power drop equation given in [24] to reckon the heterodyne crosstalk noise power (P^{HTC}_{Noise}) for MR filter Q-factor of 9000 and baud-rate (or bit-rate) of 10Gbps, which defines the detector sensitivity threshold *S* in Eq. (3). From [12], the OOK and 4-PAM modulated signals have identical frequency spectra (represented by the *sinc* function). Therefore, the power drop equation given in [24] can be equally applicable to OOK and 4-PAM signals for reckoning P^{HTC}_{Noise} by simply replacing the bit-rate in the equation with the signal baud-rate. Note that, as a single symbol in a 4-PAM signal represents two bits, the bit-rate for a 4-PAM signal is twice its baud-rate (or symbol-rate). We use Lorentzian function based equations given in [6], [28] to reckon through losses of modulator and detector MRs (P^{MMR}_{Loss} and P^{DMR}_{Loss} respectively), all of which depend on N_λ. As can be implied from [27] and [24], for a given photonic link, P^{MMR}_{Loss}, P^{DMR}_{Loss}, and P^{HTC}_{Noise} increase with increase in N_λ due to corresponding decrease in channel gap. Hence, P^{dB}_{Loss}, *S*, and therefore, P^{dB}_{Budget} depend on N_λ (Eq. (2)-(4)). Now, as evident from Eq. (2), a value of N_λ can be calculated from values of P^{dB}_{Loss} and P^{dB}_{Budget}. However, the values of P^{dB}_{Loss} and P^{dB}_{Budget} also depend on N_λ (Eq. (2)-(4)). Therefore, it can be concluded that N_λ has cyclic dependency on parameters P^{dB}_{Loss} and P^{dB}_{Budget}, i.e., N_λ depends on and also controls the parameters P^{dB}_{Loss} and P^{dB}_{Budget}. Due to this cyclic dependency, the optimal value of N_λ cannot be obtained directly from the parameters P^{dB}_{Loss} and P^{dB}_{Budget}. Therefore, we employ a search heuristic that finds the optimal value of N_λ that satisfies the constraint for given sets of input parameters, from a set of its allowable values.

$$P^{dB}_{Budget} \geq P^{dB}_{Loss} + 10\log_{10}(N_\lambda), \quad (1)$$

$$P^{dB}_{Budget} = P_{Max} - S, \quad (2)$$

$$S = 0.5\times SNR_{Target} + P^{HTC}_{Noise} + P^{Thermal}_{Noise}, \quad (3)$$

$$P^{dB}_{Loss} = P^{MMR}_{Loss} + P^{DMR}_{Loss} + P^{WGP}_{Loss} + P^{WGB}_{Loss} + P^{SpC}_{Loss} + P^{INTRF}_{Loss}, \quad (4)$$

Table 1: Definitions and values of various link design parameters

Parameter	Definition		Value
P_{Max}	Maximum allowable optical power per WG		20dBm [10]
$P^{Thermal}_{Noise}$	Thermal noise power for detector		-22dBm [20]
SNR_{Target}	Target SNR value	OOK	21.7dB [12]
		4-PAM	31.3dB [12]
P^{WGB}_{Loss}	Waveguide bending loss (dB per 90°)		0.005 [6]
P^{WGP}_{Loss}	Waveguide propagation loss		0.27dB/cm [6]
P^{SpC}_{Loss}	Splitter + coupling loss		1.2dB [6][7]
-	Laser efficiency		30% [7]
-	Detector responsivity		1.1 A/W [7]
-	Detector bandwidth		5GHz
P^{INTRF}_{Loss}	Signal interference loss	*4-PAM*-SS	4.8dB
		4-PAM-P and OOK	0dB

4.1 Search Heuristic Based Optimization

Our proposed search heuristic performs exhaustive search to find the optimal constraint-satisfying value of N_λ. To limit the cost and complexity of the comb-generating laser source [5], and to be consistent with the prior works on 4-PAM optical signaling [4] and [7], we limit the maximum allowable value of N_λ to 128. Moreover, as the flit-size of a PNoC is directly proportional to the value of N_λ, and as the flit-size is usually a power-of-two value, the allowable values of N_λ should also be power-of-two values. Because of these reasons, we give a set Λ = {128, 64, 32, 16, 8, 4, 2, 1}, which is a set of all allowable values of N_λ, as an input to our search heuristic. Based on the constraint in Eq. (1), we create an error function $ef(N_\lambda) = \{P^{dB}_{Budget} - P^{dB}_{Loss} - 10\log_{10}(N_\lambda)\}$. Then, for each

element N_λ of the set Λ, we evaluate an error value $\epsilon = ef(N_\lambda)$ using Eq. (1)-(4) and parameter values from Table 1, and create a set E of all ϵ values. For that, we evaluate P_{Noise}^{HTC}, P_{Loss}^{MMR} and P_{Loss}^{DMR} values using equations given in [24] and [6] (as mentioned earlier). All N_λ values corresponding to the positive ϵ values in set E satisfy the constraint given in Eq. (1). But we choose the N_λ corresponding to the minimum positive value ϵ_{min} from set E as the optimal value, because such N_λ is the maximum constraint-satisfying value of the number of DWDM λs.

Note that this search heuristic is equally applicable to OOK, 4-PAM-SS, and *4-PAM-P* methods. However, the optimal values of N_λ would differ for different methods, as the link design parameters such as signal-interference loss P_{Loss}^{INTRF} and SNR_{Target} differ between different methods (see Table 1).

4.2 Design for Reliability and Bandwidth

We can use the search heuristic given in Section IV-A to find the constraint-satisfying optimal value of N_λ that can achieve either maximum bandwidth in terms of aggregate bit transfer rate or desired reliability in terms of BER for designed photonic links. As can be implied from [10] and [30], traditionally, the designs of photonic links are optimized to achieve maximum bandwidth, and while doing so the reliability of the designed photonic links is usually disregarded. Therefore, in the link-bandwidth maximizing optimization frameworks presented in [10] and [30], only the noise-limited detector sensitivity is utilized, and the parameter SNR_{Target} is ignored. From Eq. (3), when the parameter SNR_{Target} is ignored to achieve maximum bandwidth, the detector sensitivity or the minimum detectable signal power (S) is set to be equal to the total noise power. As a result, the available power budget can accommodate greater number of λ-signals (corresponding to a larger value of N_λ) for a given P_{Loss}^{dB}, but each of these wavelength signals yields such a small value of signal power that can be hardly distinguished from the noise power. This results in poor SNR, BER, and communication reliability for the bandwidth-maximized photonic links.

In contrast, we utilize the parameter SNR_{Target} in Eq. (3) and the search heuristic to optimize the links for desired reliability. Introducing the parameter SNR_{Target} in our optimization framework sets the minimum detectable signal power (S) to a higher level, because of which the available power budget can accommodate only a small number of λ-signals. Nevertheless, doing so ensures that all the supported λ-signals achieve the target SNR. As a result, desired BER and communication reliability can be achieved for the reliability-optimized links. As discussed in [12], BER is a more appropriate measure of reliability than SNR, and a BER of 10^{-9} is a standard target for reliable on-chip communication. Therefore, we choose appropriate values of SNR_{Target} that correspond to BER of 10^{-9} for different signaling methods. From the equations given in [12], and as shown in Table 1, to achieve a BER of 10^{-9}, OOK and 4-PAM signaling require SNR of 21.7dB and 31.3dB, respectively. F*or given noise power,* 4-PAM *signaling requires greater SNR to achieve a BER of 10^{-9}, as a given amount of noise power impacts a* 4-PAM *signal more than an OOK signal, due to the decreased gap between different optical transmission levels of the 4-PAM signal.*

As mentioned in Section I, our goal in this paper is to evaluate the impact of the optimized designs of OOK, 4-PAM-SS, and *4-PAM-P* based photonic links on the performance, reliability and energy-efficiency of a well-known PNoC architecture: an 8-ary 3-stage CLOS PNoC [4]. To achieve this goal, we first optimize OOK, 4-PAM-SS, and *4-PAM-P* based single-waveguide photonic links for reliability to achieve BER of 10^{-9}, using our search heuristic given in Section IV-A. We optimize 4.5cm long single-waveguide links, as according to our geometric analysis of the CLOS PNoC, and from [4], the longest link of the CLOS PNoC is 4.5cm long. Optimizing a 4.5cm long single-waveguide link corresponds to finding the optimal value of N_λ using the search heuristic when all other parameters in Eq. (1)-(4) and Table 1 are set based on the length and geometrical parameters (e.g., number and degree of bends, number of splitters etc.) of the link. We set the number of DWDM λs for all the links in the CLOS PNoC to be equal to the optimal N_λ obtained for the reliability-optimized 4.5cm long link.

Table 2 gives optimal N_λ and worst-case optical loss (optical loss for the longest 4.5cm link) values for OOK, 4-PAM-SS, and *4-PAM-P* based variants of the CLOS PNoC. We direct the readers to Section V for more information on the architecture of the CLOS PNoC. As our proposed *4-PAM-P* method does not require any additional photonic structures compared to the conventional OOK method and due to the absence of interference induced signal loss, the *4-PAM-P* based variants have worst-case signal loss values that are similar to the OOK based variants (Table 2). As evident from Table 2, the 4-PAM signaling based reliability-optimized variants of CLOS PNoC (CLOS-4PAM-SS and *CLOS-4PAM-P*) render smaller N_λ than the OOK signaling based CLOS-OOK PNoC. This is because the SNR_{Target} value of 31.3dB for 4-PAM signaling is greater than the SNR_{Target} of 21.7dB for OOK signaling, which reduces the available power budget for the 4-PAM signaling based links, rendering smaller N_λ for 4-PAM signaling based variants of CLOS PNoCs. Nevertheless, as will be clear in Section V, despite having inferior bandwidth owing to the smaller N_λ, *CLOS-4PAM-P* PNoC achieves better energy-efficiency than CLOS-OOK.

Table 2: DWDM degree (optimal N_λ), optical loss and photonic area for different variants of CLOS PNoC

Configuration	Waveguide DWDM (optimal N_λ)	Worst-case optical loss (in dB)	Optical loss + 10log(N_λ) (in dB)	Photonic area (in mm^2)
Reliability-optimized PNoCs				
CLOS-OOK	64	-1.7	-19.70	2.64
CLOS-4PAM-SS	4	-6.4	-12.40	2.13
CLOS-4PAM-P	16	-1.7	-13.70	2.22
Bandwidth-neutral PNoCs				
CLOS-OOK-BN	64	-1.7	-19.70	2.64
CLOS-4PAM-SS-BN	32	-6.4	-21.40	2.50
CLOS-4PAM-P-BN	32	-1.7	-16.70	2.36

Based on the standard dimensions and sizes of photonic devices given in [10], we evaluated the total photonic area consumption of the reliability-optimized variants of CLOS PNoC. The result of this evaluation is also given in Table 2. As evident from the optimal N_λ, worst-case loss, and photonic area results of the reliability-optimized CLOS PNoCs given in Table 2, CLOS-OOK achieves the greatest bandwidth corresponding to the largest N_λ, whereas CLOS-4PAM-SS and *CLOS-4PAM-P* achieve the best values of photonic area and worst-case loss respectively. Thus, a clear winner from the CLOS-4PAM-SS, *CLOS-4PAM-P*, and CLOS-OOK PNoCs cannot be decided by looking at the reliability-optimized results given in Table 2.

To have a fair comparison and to decide the superior method out of the three signaling methods, we evaluate bandwidth neutral designs (with equal aggregate bit transfer rates) of OOK, 4-PAM-SS, and *4-PAM-P* based variants of the CLOS PNoC, referred to as CLOS-OOK-BN, CLOS-4PAM-SS-BN, and *CLOS-4PAM-P-BN*, respectively. We evaluate the worst-case loss, optimal N_λ, and

photonic area values for all three bandwidth-neutral CLOS PNoCs, which are listed in Table 2. Note that N_λ=32 for a 4-PAM signal and N_λ=64 for an OOK signal both achieve equal bandwidth (aggregate bit transfer rate), as a 4-PAM signal has 2× bit-rate than an OOK signal (Section II). As evident from Table 2, among all three bandwidth-neutral variants of CLOS PNoC, *CLOS-4PAM-P-BN* achieves the best values of worst-case optical loss (which determines required laser power) and photonic area. Therefore, it can be concluded that *4-PAM-P* method is more area- and energy-efficient than 4-PAM-SS and OOK methods.

For a fairer and more comprehensive comparison of different signaling methods, it is important to evaluate the bandwidth (aggregate bit transfer rate), reliability, and energy-efficiency of all the variants of the CLOS PNoCs listed in Table 2 for real benchmark applications and in the presence of variations. Such an evaluation is presented in the next section.

5 EVALUATION

5.1 Evaluation Setup

We considered a PNoC with an 8-ary 3-stage CLOS topology [4] for a 256-core system, with 8 clusters (C1-C8) and 32 cores in each cluster. Within each cluster, a group of four cores is connected to a concentrator. Thus, each cluster has 8 concentrators and these concentrators are connected electrically through a router for inter-concentrator communication. The CLOS PNoC uses point-to-point photonic links for inter-cluster communication, with a total of 56 waveguides being used to connect all 8 clusters of the CLOS PNoC. Each point-to-point photonic link uses either forward or backward propagating λs depending on the physical location of the source and destination clusters. This PNoC uses two laser sources to enable forward and backward communication. To power the 56 waveguides, the PNoC employs a series of 1X2, 1X7, and 1X4 splitters. Based on geometric analysis, we estimated the maximum length of a WG in the CLOS PNoC to be 4.5cm. This 4.5cm long WG act as a point-to-point link between cluster C6 and C1.

Modeling of Process Variations of MRs in CLOS PNoC: We adapt the VARIUS tool [31] to model die-to-die (D2D) as well as within-die (WID) process variations in MRs for the CLOS PNoC. We consider a 256-core chip with die size 400mm^2 at a 22nm process node. For the VARIUS tool, we use the parameters and procedures given in [6] and [14] to generate 100 process variation (PV) maps, each containing over 1 million points indicating the PV-induced resonance shift of MRs. The total number of points picked from these maps equal the number of MRs in the CLOS PNoC.

Simulation Setup: We performed benchmark-driven simulation-based analysis to evaluate the impact of various signaling methods on the performance and efficiency of the CLOS PNoC architecture. We modeled and simulated all the variants (reliability-optimized as well as bandwidth-neutral) of the CLOS PNoC given in Table 2 using a cycle-accurate NoC simulator. We evaluated performance for a 256-core single-chip architecture at a 22nm CMOS node. We kept the number of WGs and basic floorplan of the architectures constant across all the variants. We used real-world traffic from applications in the PARSEC benchmark suite [15]. GEM5 full-system simulation [16] of parallelized PARSEC applications was used to generate traces that were fed into our cycle-accurate NoC simulator. In GEM5 simulations, we set a "warm-up" period of 100M instructions and then captured traces for the subsequent 1B instructions. In our benchmark-driven simulations we evaluated total power, average latency, and energy-per-bit (EPB) values for different variants of CLOS PNoC. The results of this evaluation are given in Section V-B and Section V-C.

Moreover, to evaluate dynamic energy consumption values, we extended the photonic model of the DSENT tool [17] to include all three signaling methods (OOK, 4-PAM-SS, and *4-PAM-P*). Table 3 gives SNR_{Target} values, dynamic energy values, channel gap between adjacent λs, and detector sensitivity (S) values (evaluated using Eq. (3)) for all the variants of the CLOS PNoC considered in our evaluation. In the table, the detector sensitivity values directly correspond to SNR_{Target} values according to Eq. (3). The channel gap values in Table 3 correspond to a free-spectral range of 50nm and N_λ values given in Table 2.

Table 3: SNR_{Target}, detector sensitivity, channel gap (CG) between adjacent λs, and dynamic energy for different variants of CLOS PNoC

Reliability-optimized/Bandwidth-neutral PNoCs	SNR_{Target} (in dB)	Detector sensitivity S - in dBm	CG (in nm)	Dynamic energy (in fJ/bit)
CLOS-OOK / CLOS-OOK-BN	21.7/0	-9.2/-22	0.83	2.2
CLOS-4PAM-SS/ CLOS-4PAM-SS-BN	31.3/0	-4.4/-22	17.33/ 1.67	4.4
				2.2
				2.2
				0
CLOS-4PAM-P/ CLOS-4PAM-P-BN	31.3/0	-4.4/-22	3.47/ 1.67	2.2
				0.5
				0.2
				0

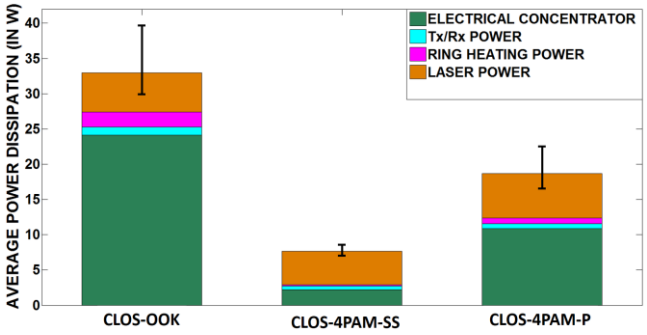

Fig. 3. Average total power dissipation comparison for different reliability-optimized configurations of the CLOS PNoC architecture.

5.2 Results for Reliability-Optimized CLOS PNoCs

Fig. 3 presents a comparison of total power dissipation values for the CLOS-OOK, CLOS-4PAM-SS, and *CLOS-4PAM-P* PNoCs. The power dissipation values in this figure are averaged across different PARSEC benchmark applications. The error bars in the figure represent maximum and minimum total power values across the 100 PV maps. As evident, compared to CLOS-OOK, *CLOS-4PAM-P* and CLOS-4PAM-SS dissipate 45.2% and 77.2% lower total power respectively. From Table 2, CLOS-OOK has the largest N_λ, whereas CLOS-4PAM-SS has the smallest N_λ. The largest value of N_λ results in the largest flit-size and hence the largest buffer-size, which in turn results in the highest power dissipation in electrical concentrators. Moreover, the largest N_λ also results in the highest number of MRs, which translates into the highest amount of MR heating power. Due to these reasons, CLOS-OOK dissipates the highest power compared to the other two variants. In contrast, the smallest value of N_λ results in the lowest power dissipation for CLOS-4PAM-SS. Moreover, from Table 2, as the reliability-optimized CLOS-4PAM-P has greater N_λ than CLOS-4PAM-SS, CLOS-4PAM-P has higher power dissipation.

Fig. 4(a), (b) present average packet latency and aggregate energy-per-bit (EPB) for all three variants of the CLOS PNoC across 12 multi-threaded PARSEC benchmarks. The error bars in Fig. 4(b) represent maximum and minimum EPB values across the 100 PV maps.

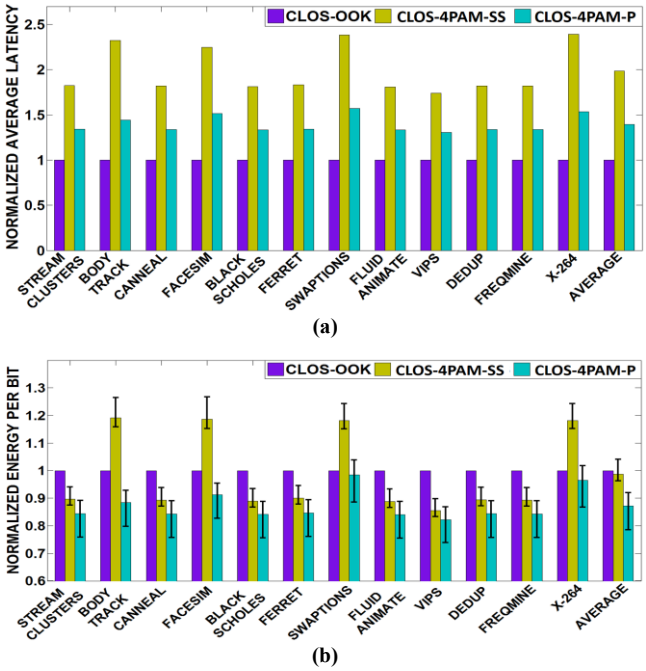

Fig. 4. (a) Average packet latency, and (b) energy-per-bit comparison for different reliability-optimized variants of CLOS PNoC across PARSEC benchmarks. All results are normalized to the baseline CLOS-OOK PNoC results.

As evident from Fig. 4(a), *CLOS-4PAM-P* achieves 29.8% lower average latency than CLOS-4PAM-SS, whereas CLOS-OOK achieves 50% and 28.8% lower average latency than CLOS-4PAM-SS and *CLOS-4PAM-P* respectively. The larger value of N_λ results in increased simultaneous bit transfers, which in turn renders lower average latency for *CLOS-4PAM-P* compared to CLOS-4PAM-SS. Similarly, the largest N_λ results in the lowest average latency for CLOS-OOK. From Fig. 4(b), *CLOS-4PAM-P* has 12.7% and 11.5% lower EPB compared to CLOS-OOK and CLOS-4PAM-SS respectively. The 4-PAM signaling used in *CLOS-4PAM-P* makes better use of MR heating power and electrical concentrator power by modulating two bits using only one modulator MR. As a result, *CLOS-4PAM-P* renders better energy-efficiency in terms of lower EPB than CLOS-OOK. Interestingly, CLOS-4PAM-SS also uses 4-PAM signaling, but the higher value of P_{Loss}^{INTRF} for 4-PAM-SS increases the signal loss, which results in very small N_λ, and hence, worse energy-efficiency (EPB) for CLOS-4PAM-SS.

As evident from Fig. 4(b), CLOS-4PAM-SS yields lower EPB than CLOS-OOK on average, as CLOS-4PAM-SS dissipates 45.2% less power, but it yields 2× average latency than CLOS-OOK on average. But for some applications such as *Bodytrack*, *Facesim*, *Swaptions*, and *X-264*, CLOS-4PAM-SS yields greater EPB than CLOS-OOK. This is because for these applications, CLOS-4PAM-SS yields greater than 2× latency compared to CLOS-OOK, the effect of which translates into greater EPB.

Note that the BER for all the three reliability-optimized variants of the CLOS PNoC is 10^{-9}. *It can be observed from these results that 4-PAM-P signaling method is more energy-efficient while being equally reliable compared to the OOK and 4-PAM-SS methods*.

5.3 Results for Bandwidth-Neutral CLOS PNoCs

As explained in Section IV-A, the bandwidths of all the bandwidth-neutral variants (CLOS-OOK-BN, CLOS-4PAM-SS-BN, and *CLOS-4PAM-P-BN*) are equal. Therefore, we do not present the bandwidth (aggregate bit transfer rate) or average latency results in this section. Instead, we only present total power dissipation, EPB, and SNR/BER results.

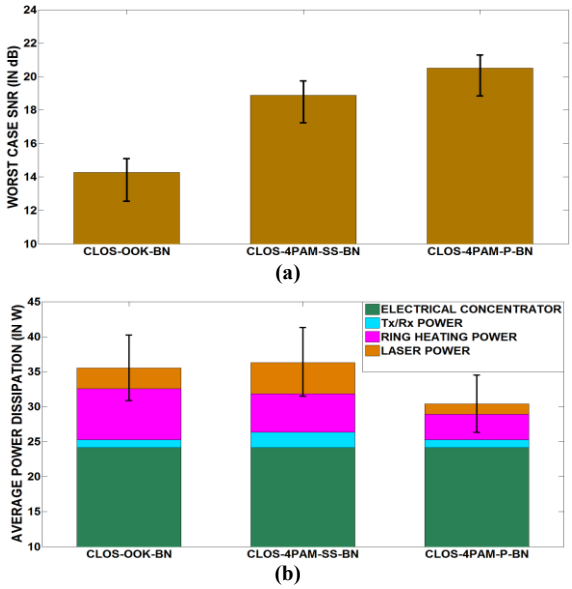

Fig. 5. (a) Worst-case SNR, (b) average total power dissipation comparison for different bandwidth-neutral configurations of the CLOS PNoC.

Fig. 5(a) plots the worst-case SNR values for different bandwidth-neutral variants of the CLOS PNoC. CLOS-OOK-BN, CLOS-4PAM-SS-BN, and *CLOS-4PAM-P-BN* yield SNR values of 14dB, 19dB and 20.5dB respectively, which translates into BER values of 6.2×10^{-3}, 1.7×10^{-2}, and 4.1×10^{-3} respectively (using the equations given in [12]). It can be implied from these results of SNR and BER that *CLOS-4PAM-P-BN* achieves greater communication reliability (corresponding to smaller BER) than CLOS-4PAM-SS-BN and CLOS-OOK-BN. As shown in Table 2 and Table 3, among all the bandwidth-neutral PNoCs, *CLOS-4PAM-P-BN* has the least signal loss and the largest channel gap. The largest channel gap results in the least value of heterodyne crosstalk noise power P_{Noise}^{HTC} [6][24]. Due to the combined effects of these factors, *CLOS-4PAM-P-BN* achieves the best SNR, BER, and communication reliability.

Fig. 5(b) gives average power dissipation for different bandwidth-neutral variants of the CLOS PNoC. *CLOS-4PAM-P-BN* has 16.9% and 19.5% lower total power dissipation compared to CLOS-OOK-BN and CLOS-4PAM-SS-BN respectively. Decrease in laser power (due to decrease in through losses), ring heating power (due to decrease in number of MRs) and dynamic energy (as shown in Table 3) contributes to decrease in total power dissipation of *CLOS-4PAM-P-BN* compared to CLOS-4PAM-SS-BN and CLOS-OOK-BN.

Lastly, Fig. 6 presents EPB values for all three bandwidth-neutral variants of the CLOS PNoC across PARSEC benchmarks. It can be seen that on an average, the *CLOS-4PAM-P-BN* has 14.6% and 16.3% lower EPB compared to CLOS-OOK-BN and CLOS-

4PAM-SS-BN respectively. Decrease in total energy consumption of *CLOS-4PAM-P-BN* compared to CLOS-OOK-BN and CLOS-4PAM-SS-BN reduces its EPB. Note that, in Fig. 5(a), (b), and Fig. 6, the error bars represent maximum and minimum across the 100 PV maps of the SNR, total power, and EPB values, respectively.

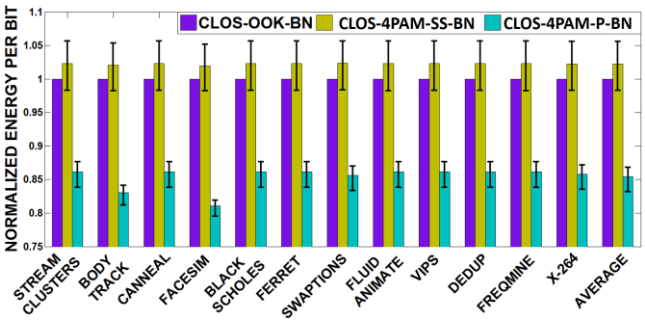

Fig. 6. Energy-per-bit comparison for different bandwidth-neutral variants of CLOS PNoC across PARSEC benchmarks. All results are normalized to the baseline CLOS-OOK-BN PNoC results.

In summary, we showed that PNoCs that utilize our proposed 4-PAM-P signaling based photonic links have greater reliability, energy-efficiency, and area-efficiency for the same bandwidth compared to the PNoCs that are comprised of 4-PAM-SS or OOK signaling based photonic links. These results corroborate the excellent capabilities of our proposed *4-PAM-P* optical signaling method in achieving high-bandwidth data transfers in PNoCs with greater reliability, area-efficiency, and energy-efficiency.

6 CONCLUSIONS

This paper presents a novel method, called *4-PAM-P*, for generating 4-PAM optical signals in PNoCs, which can double the aggregate bandwidth without increasing utilized wavelengths, photonic hardware, and incurred noise, thereby improving the bit-error-rate (BER), area-efficiency, and energy-efficiency of PNoCs. Our analysis shows that our *4-PAM-P* method achieves equal bandwidth with 4.2× better BER, 19.5% lower power, 16.3% lower energy-per-bit, and 5.6% less photonic area compared to the best known 4-PAM optical signaling method (4-PAM-SS) from prior work. Moreover, our *4-PAM-P* method achieves equal bandwidth with 1.5× better BER, 16.9% lower power, 14.6% lower EPB, and 10.6% less photonic area compared to the conventional OOK signaling method. These results corroborate the excellent capabilities of our proposed *4-PAM-P* method in achieving high-bandwidth data transfers in PNoCs with greater reliability, area- and energy-efficiency.

ACKNOWLEDGMENT

This research is supported by grants from SRC, NSF (CCF-1252500, CCF-1302693), and AFOSR (FA9550-13-1-0110).

REFERENCES

[1] L. Zhou, A. K. Kodi, “PROBE: Prediction-based optical bandwidth scaling for energy-efficient NoCs,” in NOCS, 2013.
[2] S. Pasricha, S. Bahirat, “OPAL: A Multi-Layer Hybrid Photonic NoC for 3D ICs”, in IEEE/ACM ASPDAC, Jan 2011.
[3] S. Bahirat, S. Pasricha, “METEOR: Hybrid Photonic Ring-Mesh Network-on-Chip for Multicore Architectures”, in TECS, 13(3), 2014.
[4] T. Kao et al., “Design of High Bandwidth Photonic Network-on-Chip Architectures Using Optical Multilevel Signaling,” in NOCS, 2016.
[5] C. H. Chen et al., “A comb laser-driven DWDM silicon photonic transmitter based on microring modulators”, in OE, 23(16), 2015.
[6] S.V.R. Chittamuru et al., “PICO: Mitigating Heterodyne Crosstalk Due to Process Variations and Intermodulation Effects in Photonic NoCs”, in DAC, 2016.
[7] T. J. Kao et al., “Optical Multilevel Signaling for High Bandwidth and Power-Efficient On-Chip Interconnects”, in PTL, 27(19), 2015.
[8] S. V. R. Chittamuru, et al., "SWIFTNoC: A Reconfigurable Silicon-Photonic Network with Multicast Enabled Channel Sharing for Multicore Architectures," in ACM Journal on Emerging Technologies in Computing Systems (JETC), vol. 13(4), no. 58, Feb. 2017.
[9] S. V. R. Chittamuru et al., “Crosstalk Mitigation for High-Radix and Low-Diameter Photonic NoC Architectures”, in IEEE Design and Test, 2015.
[10] I. Thakkar et al., “A Comparative Analysis of Front-End and Back-End Compatible Silicon Photonic On-Chip Interconnects”, SLIP, 2016.
[11] I. Thakkar, et al., "Mitigation of Homodyne Crosstalk Noise in Silicon Photonic NoC Architectures with Tunable Decoupling," in ACM/IEEE International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS), Oct. 2016.
[12] T. Y. Elganimi, “Performance Comparison between OOK, PPM and PAM Modulation Schemes for Free Space Optical (FSO) Communication Systems: Analytical Study”, in IJCA, 79(11), 2013.
[13] I. Thakkar, et al., "Run-Time Laser Power Management in Photonic NoCs with On-Chip Semiconductor Optical Amplifiers," in IEEE/ACM International Symposium on Networks-on-Chip (NOCS), Aug. 2016.
[14] S. K. Selvaraja., “Wafer-Scale Fabrication Technology for Silicon Photonic Integrated Circuits,” PhD thesis, Ghent University, 2011.
[15] C. Bienia et al., “The PARSEC Benchmark Suite: Characterization and Architectural Implications,” in *PACT*, 2008.
[16] N. Binkert et al., “The Gem5 Simulator,” in *Comp. Arch. News*, 2011.
[17] C. Sun et al., “DSENT—A tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling,” in NOCS, 2012.
[18] K. Padmaraju et al., “Wavelength Locking and Thermally Stabilizing Microring Resonators Using Dithering Signals”, JLT, 32(3), 2014.
[19] R. D.-Demers et al., “Ultrafast pulse-amplitude modulation with a femtojoule silicon photonic modulator”, Optica, 3(6), 2016.
[20] K. Szczerba et al., “4-PAM for High-Speed Short-Range Optical Communications”, JOCN, 4(11), 2012.
[21] A. Roshan-Zamir et al., “A 40 Gb/s PAM4 silicon microring resonator modulator transmitter in 65nm CMOS”, in OIC, 2016.
[22] S. Moazeni and V. Stojanovic. Technical Report No. UCB/EECS-2017-23. http://www2.eecs.berkeley.edu/Pubs/TechRpts/2017/EECS-2017-23.html.
[23] K. Padmaraju and K. Bergman. "Resolving the thermal challenges for silicon microring resonator devices," Nanophotonics, 3(4-5), 2014.
[24] M. Bahadori, “Optimization of Microring-based Filters for Dense WDM Silicon Photonic Interconnects”, IEEE OI, 2015.
[25] I. Thakkar, S. Pasricha, "3D-Wiz: A Novel High Bandwidth, Optically Interfaced 3D DRAM Architecture with Reduced Random Access Time," ICCD, Oct 2014.
[26] I. Thakkar, S. Pasricha, "3D-ProWiz: An Energy-Efficient and Optically-Interfaced 3D DRAM Architecture with Reduced Data Access Overhead", IEEE TMSCS, 2015.
[27] M. Bahadori et al., “Crosstalk Penalty in Microring-Based Silicon Photonic Interconnect Systems”, JLT, 34(17), pp. 4043-4052, 2016.
[28] S. V. R. Chittamuru et al., “Analyzing Voltage Bias and Temperature Induced Aging Effects in Photonic Interconnects for Manycore Computing”, SLIP, 2017.
[29] R. Wu et al., “Compact Modeling and System Implications of Microring Modulators in Nanophotonic Interconnects”, in SLIP, 2015.
[30] R. Hendry et al., "Physical layer analysis and modeling of silicon photonic WDM bus architectures," in Proc. HiPEAC Workshop, 2014.
[31] S. Sarangi et al., “Varius: A model of process variation and resulting timing errors for microarchitects,” IEEE TSM, 21(1), pp. 3 –13, 2008.
[32] Y. Xu, S. Pasricha, “Silicon Nanophotonics for Future Multicore Architectures: Opportunities and Challenges”, IEEE D&T, pp. 9-17, Sept/Oct 2014.