Crosstalk Mitigation for High-Radix and Low-Diameter Photonic NoC Architectures

Sai Vineel Reddy Chittamuru and Sudeep Pasricha
Colorado State University

Editor's notes:
Photonic Network-on-chip (PNoC) is a promising alternative to design low-power and high-bandwidth interconnection infrastructure for multicore chips. The micro ring resonators, which are essential building blocks for designing PNoCs are susceptible to crosstalk that can notably degrade signal-to-noise ratio (SNR), reducing reliability of PNoCs. This paper proposes two novel encoding mechanisms to improve worst-case SNR by reducing crosstalk noise in microring resonators used within high-radix and low-diameter crossbar-based PNoCs.
—Partha Pratim Pande, Washington State University

With several hundreds of on-chip cores expected to become a reality in the near future, electrical networks-on-chip (ENoCs) are projected to suffer from cripplingly high-power dissipation and limited performance [1]. Recent developments in silicon nanophotonics have enabled the integration of on-chip photonic interconnects with CMOS circuits, ushering in photonic networks-on-chip (PNoCs) that can offer ultra-high bandwidth, reduced power, and lower latency than ENoCs.

Several PNoC architectures have been proposed with a crossbar topology (e.g., [2], [3]). These architectures are built using silicon photonic devices such as microring-resonators (MRs) and silicon-waveguides, and employ dense-wavelength-division-multiplexing (DWDM) for bit-parallel and packet-serial type of transmission, where a large number of wavelengths are multiplexed in a waveguide to enable high-bandwidth parallel data transfers. Unfortunately, MRs suffers from intrinsic crosstalk-noise and power-loss due to their design imperfections. The crosstalk noise severely impacts PNoCs, especially crossbar architectures with high MR counts, where the generated crosstalk is intensified, leading to transmission errors. For example, the Corona [2] crossbar architecture has worst-case SXR of 14 dB [10] in its data channels, which is insufficient for reliable data communication, as its corresponding bit-error-rates (BER) are very high, in the order of 10^{-3}.

Crosstalk in DWDM-based PNoCs mainly occurs due to inefficient coupling in ring-detectors, with nonresonant-wavelengths closer to the detector resonance-wavelengths creating greater crosstalk-noise. In the electrical domain, crosstalk occurs when adjacent wires simultaneously transition in opposite directions. The code-words used in the electrical domain are not directly applicable to the photonic domain. For example, forward-error-correcting (FEC) codes that are effective in correcting erroneous bit-flips in the electrical domain utilize

Color versions of one or more of the figures in this paper are available online at http://ieeexplore.ieee.org.
Digital Object Identifier 10.1109/MDAT.2015.2414417
Date of publication: 18 March 2015; date of current version: 04 June 2015.

code-words with adjacent 1's that cannot improve optical-signal-to-noise ratio (SXR) in the photonic domain. Thus different techniques to mitigate crosstalk noise and improve reliability are needed for PNoCs.

We observe that when transmitting data in PNoCs, crosstalk noise in MRs depends on the characteristics of data values propagating in the photonic waveguide. Therefore we propose two novel techniques to intelligently reduce undesirable data value occurrences in a photonic waveguide. These techniques are easily implementable in any existing DWDM-based photonic crossbar without requiring major modifications to the architectures, unlike previously proposed crosstalk mitigation techniques (e.g., [6]) that are targeted to reduce crosstalk in specific architectures by requiring modifications to their router designs. Our novel contributions in this work are the following.

- Design a crosstalk mitigation technique with 5-bit encoding (PCTM5B) to improve worst-case SXR for DWDM-based photonic crossbar PNoCs;
- Introduce another crosstalk-mitigation scheme with 6-bit encoding (PCTM6B), that more aggressively improves SXR but with relatively higher EDP overhead;
- Validate our schemes by implementing them on well-known crossbar PNoCs: Corona [2] and Firefly [3], for real-world multithreaded PARSEC [4] benchmarks.

Related work

Several prior works have performed photonic crosstalk analysis at the device-level and architecture-level. The device-level efforts analyze crosstalk behavior for single waveguide crossings (e.g., [5]) and for one or few photonic switching elements with MRs [8]. Results from these efforts show that crosstalk is very small at the device-level. However, at the architecture level, prior work (e.g., [6]) indicates that crosstalk has a significant impact on the SXR of PNoCs because of the presence of several waveguide-crossings and switching-elements. Minimum SXR was shown to be key limiting factor in the design of mesh-based PNoCs [6]. Further, [7] showed that in fat-tree-based PNoCs, crosstalk-noise power is higher than signal power when on-chip core counts exceed 128.

The above works focus on single-wavelength PNoCs, where crosstalk is generated from a single wavelength. A few prior works have also explored crosstalk in DWDM-based PNoCs where multiple wavelengths coexist in a waveguide. A cascaded MR-based modulator structure is proposed in [8] for low-density DWDM waveguides, with an extinction ratio of 13 dB and negligible crosstalk. In [9], losses in a similar multiwavelength MR-based structure are measured. Though crosstalk appears negligible in these works where only four-wavelength DWDM waveguides are considered, in crossbar PNoC architectures such as Corona [2] that use 64-wavelength DWDM, there is significant crosstalk noise. The results in [10] demonstrate the damaging impact of crosstalk-noise in Corona, where the worst-case SXR is estimated to be 14 dB in data waveguides, which is insufficient for reliable data communication. A methodology to salvage network-bandwidth loss due to process-variation-drifts was proposed in [13], which reorders microrings and trims them to nearby wavelengths. In [4], [11], reliability aware multiple-segmented-bus (MSB) based PNoCs are proposed to enable data transfers with low BER. But this work does not address crosstalk reliability issues. Other efforts focus on architecture-specific crosstalk-mitigation [6], [7] by changing the physical design of PNoC routers. However, to date, no prior work has proposed generalized approaches to improve SXR in an entire class of PNoC architectures, as we do in this work, for bit-parallel and packet-serial photonic data-transmission.

Analytical models for crosstalk analysis in DWDM-based PNoC architectures

Overview of MR operation in DWDM-based PNoCs

DWDM-based PNoC architectures utilize photonic devices such as microring-resonators (MRs), photonic waveguides, splitters, and trans-impedance-amplifiers (TIAs). MRs in particular are essential to modulate light for transmission of data at a source-node (*data-modulation-phase*). MRs also detect light-modulated data from the waveguide at the destination-node (*data-detection-phase*) and subsequently generate proportional electrical signals that are amplified by TIAs. Each source node requires optical power/signals that are made available in the PNoC via power waveguides and splitters. An unfortunate property of silicon photonic waveguides is that

signal propagation is lossy, i.e., the light signal is subject to losses such as through-loss, modulating-loss, and detecting-loss in MRs, propagation-loss and bending-loss in waveguides, and splitting-loss in splitters. Such losses negatively impact SXR in waveguides.

At any point in time in a photonic-waveguide, MRs are either in-resonance or out-of-resonance with respect to the incident wavelengths. In the resonance-mode, an MR couples light of a wavelength from the waveguide when its circumference is an integer multiple of that wavelength. Different-sized MRs in a DWDM-waveguide are thus required to simultaneously modulate data on different available wavelengths. These MRs in DWDM-based PNoCs suffer from intrinsic crosstalk-noise and power-loss. Figure 1a–d shows crosstalk noise (as dotted/dashed lines) in modulator and detector MRs during typical modulation/detection phases in the DWDM-waveguide. Whenever a modulator modulates a "0" or a detector detects a "1" from a particular wavelength by removing the light pulse, there is also crosstalk generated in the waveguide, as shown in Figure 1a and d.

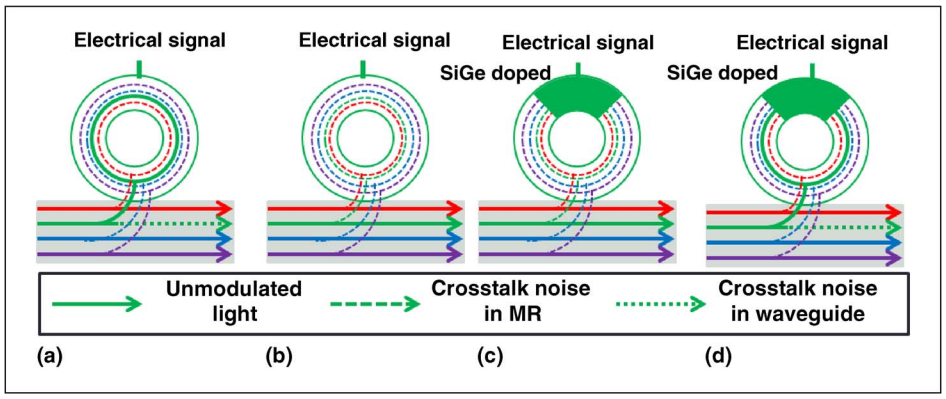

Figure 1. MR operation phases in DWDM-based waveguides (a) modulator modulating in resonance-wavelength (b) modulator in passing (through) mode (c) detector in passing-mode (d) detector in detecting-mode.

Analytical models for crosstalk-noise and signal-power

In this work we consider crosstalk in DWDM-waveguides for the Corona PNoC architecture enhanced with token-slot arbitration [2] and the Firefly PNoC architecture [3]. In DWDM-based waveguides in both architectures, data-transmission requires modulating light using a series of MR-modulators equal to the number of wavelengths supported by DWDM. Similarly, data-detection at the receiver requires a group of MR-detectors equal to the number of DWDM wavelengths. We present analytical equations that model worst-case crosstalk-noise power, maximum power-loss, and SXR in the MR-detector groups (similar equations are applicable to MR-modulator groups). We have validated these analytical models against device-level works [8], [9]. In these analytical models we assume negligible intermodulation crosstalk. In the interest of brevity, we only present models and a description for the Corona PNoC, although our evaluation show results for both the Corona and Firefly PNoCs. We refer the reader to [2], [3] for details of the photonic crossbar topology and protocols employed in the Corona and Firefly PNoC architectures.

The notations for parameters used in the analytical equations are shown in Table 1. Corona is designed for a 256 core single-chip platform, with

Table 1 Notations for photonic power-loss, crosstalk-coefficients and model-parameters [10].

Notation	Parameter type	Parameter value
L_P	Propagation loss	-0.274 dB per cm
L_B	Bending loss	-0.005 dB per 90°
L_{MI}	Inactive modulator through loss	-0.0005 dB
L_{MA}	Active modulator power loss	-0.6 dB
L_{DP}	Passing detector through loss	-0.0005 dB
L_{DD}	Detecting detector power loss	-1.6 dB
L_{S12}	1X2 splitter power loss	-0.2 dB
L_{S14}	1X4 splitter power loss	-0.2 dB
L_{S15}	1X5 splitter power loss	-0.2 dB
L_{S16}	1X6 splitter power loss	-0.2 dB
X_{MA}	Active modulator	-16 dB
X_{DD}	Detecting detector	-16 dB
Q	Q-factor of MR	9000
FSR	Free spectral range	62 nm
	Other model parameter notations	
Φ (i, j)	Coupling factor between i^{th} microring resonators and j^{th} wavelengths in waveguide	
L	Photonic path length in cm	
B	Number of bends in photonic path	
λ_j	Resonance wavelength of MR	
R_{S12}	Splitting factor for 1X2 splitter	
R_{S14}	Splitting factor for 1X4 splitter	
R_{S15}	Splitting factor for 1X5 splitter	
R_{S16}	Splitting factor for 1X6 splitter	

cores grouped into 64 clusters, and four cores-per-cluster. For intercluster communication, Corona uses a photonic-crossbar topology with 64 data-channels. Each channel consists of four multiple-write-single-read (MWSR) waveguides with 64-DWDM in each waveguide. As modulation occurs on both positive and negative edges of the clock in Corona, 512-bits (cache-line size) can be modulated and inserted on 4-MWSR waveguides in a single cycle by a sender-node. A data-channel starts at a cluster called "home-cluster," traverses other clusters (where modulators can modulate light in this channel and detectors can detect this light), and finally ends at the home-cluster again, at a set of detectors (optical termination).

A power-waveguide supplies optical power to each of the 64 data-channels at its home-cluster, through a series of 1X2 splitters starting from home-cluster 1 to 64. In each home-cluster, optical-power is distributed among 4-MWSR waveguides equally using a 1×4 splitter. As all 1×2 splitters are present before the last (64th) channel, this channel suffers the most signal-power-loss. Thus, the worst-case signal and crosstalk-noise power exists in the detector group of the 64th cluster node, and this node is defined as the worst-case power-loss-node (N_{WCPL}). For this node, signal-power [($P_{signal}(j)$)] and crosstalk-noise-power [($P_{noise}(j)$)] received at each detector j are expressed in (1) and (2) [10]

$$P_{signal}(j) = L_{DD}\Phi(j,j)P_S(j,j) \tag{1}$$

$$P_{noise}(j) = L_{DD}P_N(j,j) + \sum_{i=1}^{n} \Phi(i,j)(P_S(i,j) + P_N(i,j))(i \neq j). \tag{2}$$

The parameters in the above equations are defined below: (See equation at bottom of page.)

$P_S(i,j)$ in (3) is the signal-power of the ith-wavelength received before the jth-detector. Similarly in (5), $P_N(i,j)$ is the crosstalk-noise power of the ith-wavelength before the jth-detector. K_S and K_N in (6) and (7) represent signal and crosstalk-noise power-losses before the detector group of N_{WCPL}. $\psi(i,j)$ in (4a) represents signal power-loss of the ith-wavelength before the jth-detector within the detector group of N_{WCPL}. $\psi(i,j)$ in (4b) is the crosstalk coupling-factor of the ith-wavelength and the jth-detector. Finally, we can define SXR(j) of the jth-detector of N_{WCPL} as the ratio of $P_{signal}(j)$ to $P_{noise}(j)$, as shown in (8):

$$SXR(j) = \frac{P_{signal}(j)}{P_{noise}(j)} \tag{8}$$

$$P_S(i,j) = K_S\psi(i,j)P_{in}(i) \tag{3}$$

$$\psi(i,j) = \begin{cases} X_{DD}(L_{DP})^{j-1} & (j-1 \geq i \text{ and } D_B = 1) \\ (L_{DP})^{j-1} & (j-1 < i \text{ and } D_B = 1) \\ X_{MA}X_{DD}(L_{DP})^{j-1} & (j-1 \geq i \text{ and } D_B = 0) \\ X_{MA}(L_{DP})^{j-1} & (j-1 < i \text{ and } D_B = 0) \end{cases} \tag{4a}$$

$$\Phi(i,j) = \frac{\delta^2}{\left((i-j)\frac{FSR}{n}\right)^2 + \delta^2}, \text{ Here } \delta = \lambda_j/2Q \tag{4b}$$

$$P_N(i,j) = \begin{cases} 0, & \text{If } j > i \text{ and } D_B = 1 \\ K_N(L_{DP})^j P_{in}(i), & \text{if } j \leq i \text{ and } D_B = 1 \\ 0, & \text{If } j > i \text{ and } D_B = 0 \\ X_{MA}K_N(L_{DP})^j P_{in}(i), & \text{if } j \leq i \text{ and } D_B = 0 \end{cases} \tag{5}$$

$$K_S = \begin{cases} (R_{S14})(L_{S14})(L_P)^L(L_B)^B(L_{MI})^{64\times 63}(D_B = 1) \\ (R_{S14})(L_{S14})(L_P)^L(L_B)^B(L_{MI})^{64\times 62+63}(D_B = 0) \end{cases} \tag{6}$$

$$K_N = \begin{cases} (R_{S14})(L_{S14})(L_P)^L(L_B)^B X_{MA}(L_{MI})^{64\times 62+63}(D_B = 1) \\ (R_{S14})(L_{S14})(L_P)^L(L_B)^B X_{MA}(L_{MI})^{64\times 62+62}(D_B = 0) \end{cases} \tag{7}$$

These equations are sufficient to analyze signal and crosstalk-noise power during the detection of ones (D_B = '1') and zeros (D_B = '0') in a photonic waveguide. The next section uses these models to discuss how crosstalk mitigation techniques impact SXR.

Techniques to mitigate crosstalk noise

Crosstalk noise in the detectors of DWDM-based PNoCs is caused mainly due to inefficient coupling of MRs, as MR-detectors in their detecting mode not only couple photonic-power from their resonance-wavelengths but also couple some photonic-power from other wavelengths in the waveguide. This coupling-factor (Φ) increases with a decrease in gap between resonant and nonresonant wavelengths of an MR-detector ((4b)). Thus nonresonant wavelengths closer to the detector resonance wavelengths create greater crosstalk-noise. Additionally, crosstalk-noise increases with increase in signal power of nonresonant wavelengths.

Based on the above observations, crosstalk noise can be mitigated by placing one or more 0's adjacent to 1's in the data in the waveguide, to reduce photonic signal-strength of immediate nonresonant wavelengths (adjacent wavelengths in DWDM). In this section, we present two techniques for mitigation of crosstalk noise in DWDM-based PNoCs that utilize this mechanism. The two techniques (PCTM5B, PCTM6B) employ 5-bit and 6-bit encoding for every 4-bit data block to reduce photonic-signal-strength of the immediate nonresonant wavelengths. The area/power/delay overheads of these techniques are discussed in Section V-A.

PCTM5B encoding technique

Table 2 shows the 5-bit codes proposed in the *PCTM5B* scheme, to replace 4-bit data words. To implement this encoding technique on a 64-bit word, 16 additional bits are required, which increases the number of MRs by 25%. To facilitate simultaneous transfer of an entire packet in Corona, which requires 512-bits before encoding, we increase DWDM-degree in MWSR-waveguides from 64 to 65 and increase MWSR-waveguides in each channel from 4 to 5. To distribute optical-power between these waveguides, there is also a need to replace 1X4 splitters with 1X5 splitters. Therefore (6), (7) for worst-case signal and crosstalk-noise power are changed to (9) and (10) shown at the bottom of the page.

Table 2 Code words for encoding techniques.

Code words for PCTM5B technique			
Data Word	**Code Word**	**Data Word**	**Code Word**
0000	00000	1000	01000
0001	00001	1001	01001
0010	00010	1010	01010
0011	10101	1011	10100
0100	00100	1100	01100
0101	00101	1101	10010
0110	00110	1110	10001
0111	10110	1111	10000
Code words for PCTM6B technique			
Data Word	**Code Word**	**Data Word**	**Code Word**
0000	000000	1000	001000
0001	000001	1001	001001
0010	000010	1010	001010
0011	100000	1011	010100
0100	000100	1100	100010
0101	000101	1101	010010
0110	010101	1110	010001
0111	100001	1111	010000

PCTM6B encoding technique

The codes used in this 6-bit encoding technique are shown in Table 2. This encoding technique requires 32 additional bits for a 64-bit data word, and increases the number of MRs by 50%. To facilitate simultaneous transfer of an entire packet in Corona, which requires 512-bits before encoding, we increase DWDM-degree in MWSR-waveguides from 64 to 66 and increase MWSR-waveguides in each

$$K_S = \begin{cases} (\mathrm{R}_{S15})(L_{S15})(\mathrm{L}_P)^L(\mathrm{L}_B)^B(\mathrm{L}_{MI})^{65\times 63}(D_B = 1) \\ (\mathrm{R}_{S15})(L_{S15})(\mathrm{L}_P)^L(\mathrm{L}_B)^B(\mathrm{L}_{MI})^{65\times 62+64}(D_B = 0) \end{cases} \quad (9)$$

$$K_N = \begin{cases} (\mathrm{R}_{S15})(L_{S15})(\mathrm{L}_P)^L(\mathrm{L}_B)^B X_{MA}(\mathrm{L}_{MI})^{65\times 62+64}(D_B = 1) \\ (\mathrm{R}_{S15})(L_{S15})(\mathrm{L}_P)^L(\mathrm{L}_B)^B X_{MA}(\mathrm{L}_{MI})^{65\times 62+63}(D_B = 0) \end{cases} \quad (10)$$

channel from 4 to 6. To distribute optical power between these waveguides, there is also a need to replace 1X4 splitters with 1X6 splitters. The modified versions of (6) and (7) for worst-case signal and crosstalk-noise power in Corona are shown in (11) and (12) at the bottom of the page.

Evaluation studies

Evaluation methodology

To evaluate the proposed crosstalk-noise mitigation schemes, we implement them on two well-known crossbar-based PNoC architectures: Corona [2] and Firefly [3]. We modeled and simulated our schemes on these architectures using a cycle-accurate NoC simulator. We considered a 256-core single-chip architecture at 22 nm for performance analysis. A system-level simulation was performed with the open-source GEM5 architectural simulator with 256 ARM-based cores running parallelized PARSEC benchmarks, to generate traces that were fed into our cycle-accurate NoC simulator. We set a "warm-up" period of 100-million instructions and then captured traces for the subsequent 1-billion instructions. We performed geometric calculations for a 20 mm $\times$ 20 mm chip size, to determine lengths of MWSR and SWMR waveguides in the Corona and Firefly PNoCs, respectively. Based on this analysis, we estimated the time for light to travel from the first to the last node as 8-cycles at 5 GHz clock frequency in both PNoCs. We used a total packet size of 512-bits as advocated in these architectures, and a DWDM wavelength range in the C and L bands [9], with a starting wavelength of 1530 nm.

We increased photonic hardware of Corona and Firefly to have a minimal performance (latency, bandwidth) impact and to still enable transfer of an entire encoded-packet from source to destination in one-cycle. PCTM5B (PCTM6B) requires 25% (50%) increase in number of micro-rings and waveguides over the baseline Corona and Firefly architectures. More precisely, PCTM5B has an area overhead of 5.98 mm^2 (electrical) and 8.98 mm^2 (photonic) for Corona; and 11.95 mm^2 (electrical) and 18.69 mm^2 (photonic) for Firefly. PCTM6B has an area overhead of 9.34 mm^2 (electrical) and 17.97 mm^2 (photonic) for Corona; and 18.69 mm^2 (electrical) and 25.34 mm^2 (photonic) for Firefly. These area overheads are reasonably low compared to the much larger footprint of the chip (400 mm^2). The static and dynamic energy consumption of routers and concentrators in the Corona and Firefly PNoCs is calculated with the open-source DSENT tool. Both our schemes also have electrical power overhead: PCTM5B has 0.2 W and 0.4 W overhead, and PCTM6B has 0.6 W and 1.2 W overhead for Corona and Firefly, respectively.

We estimated electrical area and power overhead using gate-level analysis and the open-source CACTI-6.5 tool for memory/buffers. Photonic area overhead is estimated based on the physical dimensions of waveguides, MRs and splitters [9]. For energy consumption of photonic devices, we adopt parameters from [10], [12], with 0.42 pJ/bit for every modulation and detection event, and 0.18 pJ/bit for driver circuits of modulators and photodetectors. We used photonic-loss values for photonic components, as shown in Table 1, to obtain the photonic laser-power-budget and the corresponding electrical laser power. We consider a one-cycle overhead for both encoding and decoding of data in PCTM5B and PCTM6B, based on our circuit-level analysis at 5 GHz. While we consider the baseline token management scheme from [2] for Corona, more sophisticated token management schemes can be employed to further reduce delay overheads of PCTM5B and PCTM6B. Lastly, as our scope of work is limited to optical crosstalk analysis, we consider optical single to noise ratio (SXR) as a measure for reliability. A detailed analysis and optimization of the resulting SNR in the electrical domain at the optical receivers, due to factors such as thermal- and shot-noise is beyond the scope of this work.

$$K_S = \begin{cases} (\mathrm{R}_{S16})(L_{S16})(\mathrm{L}_P)^L(\mathrm{L}_B)^B(\mathrm{L}_{MI})^{66\times 63}(D_B = 1) \\ (\mathrm{R}_{S16})(L_{S16})(\mathrm{L}_P)^L(\mathrm{L}_B)^B(\mathrm{L}_{MI})^{66\times 62+65}(D_B = 0) \end{cases} \tag{11}$$

$$K_N = \begin{cases} (\mathrm{R}_{S16})(L_{S16})(\mathrm{L}_P)^L(\mathrm{L}_B)^B X_{MA}(\mathrm{L}_{MI})^{66\times 62+65}(D_B = 1) \\ (\mathrm{R}_{S16})(L_{S16})(\mathrm{L}_P)^L(\mathrm{L}_B)^B \mathrm{X}_{MA}(\mathrm{L}_{MI})^{66\times 62+64}(D_B = 0) \end{cases} \tag{12}$$

Evaluation results with corona architecture

Utilizing the models presented in Sections III and IV, we calculate the received crosstalk-noise and SXR at detectors for the node with worst-case power-loss (N_{WCPL}), which corresponds to detectors in cluster 64 for Corona. We compared the baseline Corona PNoC with fair token-slot arbitration [2] but without any crosstalk-enhancements, with two variants of the architecture corresponding to the two crosstalk-mitigation strategies proposed in this work. The worst-case SXR for the baseline Corona PNoC occurs when all the 64-bits of a received data word in a waveguide are 1's. However, for the implementations of Corona with our crosstalk-mitigation techniques, this is not the case, i.e., each detector in cluster 64 has a worst-case SXR for a different pattern of 1's and 0's in the received data word. We used our analytical models to determine these unique worst-case patterns for each of the techniques when used with Corona.

Figure 2a–c show detector signal power-loss, crosstalk-noise power-loss, and SXR corresponding to the detectors in the 64th cluster for the baseline and two variants of the Corona architecture. Note that the number of detectors in the node (x-axis) varies across the proposed techniques and depends on the number of data bits transmitted in the data waveguide for each technique, as discussed in Section IV. Figure 2b indicates that worst-case SXR (lowest value of the bars, which represent SXR in detectors) improves notably over the baseline shown in Figure 2a when using PCTM5B. However the improvement is on the lower side for the remaining detectors. Figure 2c shows that PCTM6B improves worst-case SXR marginally over PCTM5B, but does a better job of improving SXR significantly for most detectors.

From Figure 2, the worst-case SXR results for the baseline, PCTM5B and PCTM6B techniques are 21.74, 24.13, and 25.50, respectively. The worst-case SXR is obtained at the 42nd detector of the 64th cluster in the baseline case; whereas for PCTM5B and PCTM6B, worst-case SXR occurs at the 45th and 48th detectors of the same cluster, respectively. The worst-case SXR for Corona from [10]was shown to be close to 14 dB—our baseline result ($SXR = 21.74$) when converted to dB scale ($SXR_{db} = 13.4$ dB) is in line with those results, with a slight difference due to the use of enhanced token-slot arbitration in Corona, compared to [10]. From these results it can be surmised that Corona with PCTM5B and PCTM6B techniques has 11% and 18% improvements in worst-case SXR compared to the baseline. Both PCTM5B and PCTM6B eliminate occurrences of "111" in a data word and also limit occurrences of "11," to reduce crosstalk-noise in detectors.

Figure 3a shows results that quantify average-packet-latency and energy-delay-product (EDP) for the three Corona configurations, across twelve multithreaded PARSEC benchmarks. It can be observed that on average, Corona configurations with PCTM5B and PCTM6B have a 9% higher average-latency compared to the baseline. The additional delay due to encoding and decoding of data with PCTM5B and PCTM6B contributes to this latency increase. The Corona configurations with PCTM5B and PCTM6B have 26.5% and 46.2% higher EDP compared to the baseline, respectively. This increase in EDP is not only due to the increase in average latency, but also due to the addition of extra bits for encoding and decoding, which leads to an increase in the amount of photonic hardware in the architectures (more number of MRs, more complex splitters), which increases static energy. Dynamic energy also increases in these architectures, but to a much lesser extent.

Evaluation results with firefly architecture

To understand how our crosstalk-mitigation techniques behave when ported to a different PNoC architecture, we integrated the techniques with the Firefly [3] crossbar-based PNoC architecture. Unlike the MWSR waveguides used in Corona, Firefly employs reservation-assisted single-write-multiple-reader (R-SWMR) data-waveguides. Each data-channel in Firefly consists of 8 single-write-multiple-read (SWMR) waveguides, with 64-DWDM in each waveguide. Firefly uses only one-eighth of the MRs on each data waveguide compared to Corona, as only eight nodes are capable of accessing each SWMR waveguide. We considered a power waveguide in Firefly similar to that used in Corona and determined that the worst-case-power-loss node (N_{WCPL}) is at the detectors of C_4R_0, which is the router 0 of cluster 4 in Firefly. Similar to Corona, in Firefly the worst-case signal ($P_{signal}(j)$) and noise power ($P_{noise}(j)$) in the detectors of router C_4R_0 are calculated using (1)–(5) and SXR is calculated by (8). But as Firefly has fewer number of MRs in its data channels, this in turn changes the signal and crosstalk noise power losses before the

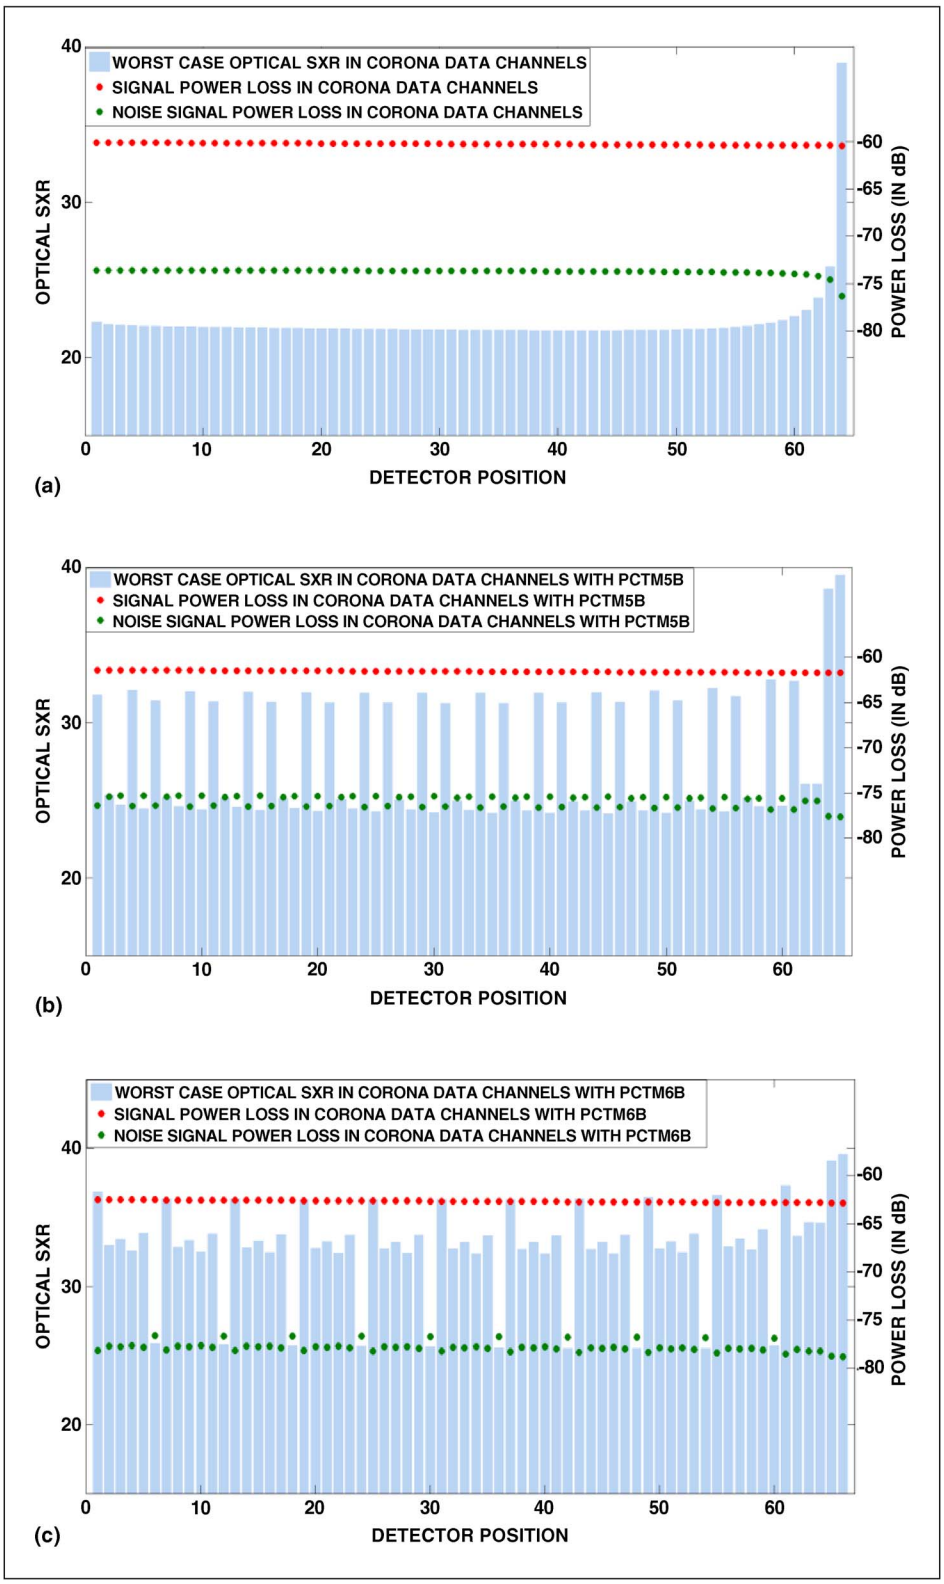

Figure 2. Detector-wise signal power-loss, crosstalk-noise power-loss, and minimum optical-SXR in worst-case power-loss node for Corona (a) baseline with 64-detectors (b) PCTM5B with 65-detectors (c) PCTM6B with 66-detectors.

detector group of N_{WCPL}, so K_S and K_N from (6) and (7) are modified to capture the architecture-specific requirements for Firefly.

To implement our PCTM5B and PCTM6B crosstalk-mitigation techniques on Firefly, we propose to make changes similar to that made for Corona, in terms of an appropriate increase in the number of waveguides in each data channel and DWDM in each waveguide, as dictated by the technique. The worst-case SXR comparison of the baseline Firefly architecture and Firefly configurations with the PCTM5B and PCTM6B techniques is shown in Figure 3b (top-figure). It can be observed that Firefly with the PCTM5B and PCTM6B techniques has 10.5% and 16.5%, improvements in worst-case SXR over the baseline configuration.

We also ran simulations with 12 applications from the PARSEC benchmark-suite and obtained normalized average packet-latency and normalized average-EDP for the Firefly configurations, just as we did for the Corona configurations. These results are shown in Figure 3b (bottom-figure), with vertical lines on the bars showing the range of values obtained across the 12 benchmarks. It can be observed that on average, Firefly with PCTM5B and PCTM6B has 9.8% higher average latency compared to the baseline configuration. The PCTM5B and PCTM6B techniques also increase EDP in Firefly by 10% and 12.8%, respectively. This increase in latency and EDP of the encoding techniques is due to the additional clock-cycles needed by the encoding and

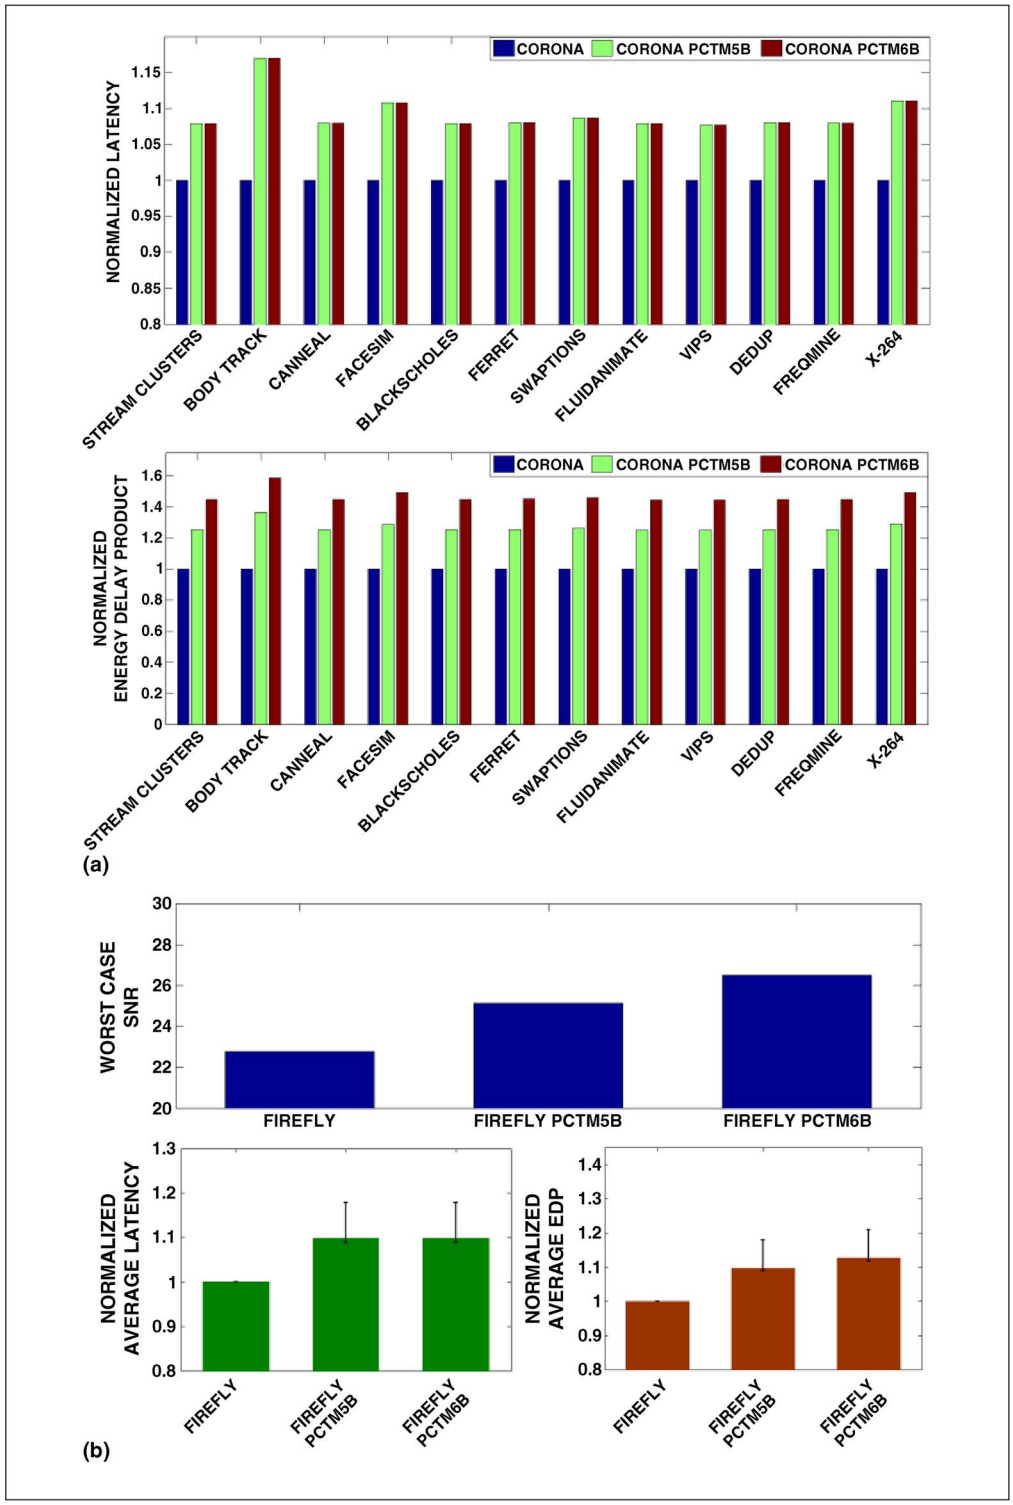

Figure 3. (a) Normalized-latency and normalized energy-delay-product (EDP) comparison between Corona baseline and Corona with PCTM5B and PCTM6B, for PARSEC benchmarks. Results are normalized to the baseline Corona results; (b) Worst-case optical-SXR (on-top), normalized average-latency (bottom-left) and EDP (bottom-right) for PARSEC benchmarks running on the baseline Firefly architecture and Firefly with PCTM5B and PCTM6B.

decoding phases as well as the additional bits used by these mechanism, that translate to greater photonic hardware which increases static/dynamic power overheads.

Summary of results and observations

From the results presented in the previous sections, we can summarize that both of our crosstalk-mitigation techniques significantly reduce crosstalk-noise and improve SXR in photonic data waveguides. Our techniques have a less than 10% latency overhead (Figure 3) and less than 5% throughput overhead (results omitted for brevity), on average. The EDP overheads of the techniques are much lower on architectures optimized for physical-layouts such as Firefly, than on nonoptimized architectures such as Corona. The PCTM5B technique is a good option to implement on DWDM-based PNoC crossbars with stringent limitations on area and energy overheads and where modest improvements to reliability are sufficient. The PCTM6B technique is a more viable choice for DWDM-based PNoC crossbars that are biased more towards reliability than energy consumption or area concerns, i.e., where higher energy and area overheads are an acceptable price to pay for greater reliability. Finally, while it is possible to extend our work to create higher order encoding scheme variants with better reliability, e.g., PCTM7B and PCTM8B, we believe that their prohibitively high power and EDP overheads limit their practical applicability.

We have presented two crosstalk mitigation techniques for the reduction of crosstalk noise in the detectors of DWDM-based PNoC architectures with crossbar topologies. These techniques (PCTM5B, PCTM6B) show interesting trade-offs between reliability, performance, and energy overhead across two different crossbar-based PNoC architectures. Our analysis on the well-known Corona and Firefly PNoCs has shown that the PCTM5B and PCTM6B techniques can notably improve worst-case SXR by up to 18%. ■

Acknowledgment

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

■ References

[1] J. D. Owens et al., "Research challenges for on-chip interconnection networks," in *Proc. IEEE Micro*, Chicago, IL, USA, 2007, pp. 96–108.

[2] D. Vantrease et al., "Light speed arbitration and flow control for nanophotonic interconnects," in *Proc. IEEE MICRO*, New York, NY, USA, 2009, pp. 304–315.

[3] Y. Pan et al., "Firefly: Illuminating future network-on-chip with nanophotonics," in *Proc. ISCA*, New York, NY, USA, 2009, pp. 429–440.

[4] P. K. Kaliraj et al., "Performance Evaluation of Reliability Aware Photonic Network on Chip Architectures," in *Proc. IGCC*, San Jose, CA, USA, 2012, pp. 1–6.

[5] C. H. Chen, "Waveguide crossings by use of multimode tapered structures," in *WOCC*, Kaohsiung, Taiwan, 2012, pp. 130–131.

[6] Y. Xie et al., "Crosstalk Noise and Bit Error Rate Analysis for Optical Network-on-Chip," in *Proc. DAC*, Anaheim, CA, USA, 2010, pp. 657–660.

[7] M. Nikdast et al., "Fat-tree-based optical interconnection networks under crosstalk noise constraint," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, pp. 156–169, 2014.

[8] Q. Xu et al., "Cascaded silicon micro-ring modulators for wdm optical interconnection," *Opt. Exp.*, vol. 14, vol. 20, pp. 9431–9435, 2006.

[9] S. Xiao et al., "Modeling and measurement of losses in silicon-on-insulator resonators and bends," *Opt. Exp.*, vol. 15, no. 17, pp. 10 553–10 561, 2007.

[10] L. H. K. Duong et al., "A case study of signal-to-noise ratio in ring-based optical networks-on-chip," *IEEE Des. Test*, pp. 55–65, 2014.

[11] I. Datta et al., "Design methodology for optical interconnect topologies in NoCs with BER and transmit power constraints," *J. Lightwave Technol.,* vol. 32, no. 1, pp. 163–175, 2014.

[12] P. Grani and S. Bartolini, "Design options for optical ring interconnect in future client devices," *J. Emerging Technol. Comput. Syst.*, vol. 10, no. 4, 2014, Art. ID. 30.

[13] Y. Xu et al., "Tolerating process variations in nanophotonic on-chip networks," in *Proc. ISCA*, Portland, OR, USA, 2012, pp. 142–152.

Sai Vineel Reddy Chittamuru received the BTech degree in electrical engineering from Indian Institute of Technology Kharagpur (IIT-KGP), India, in 2011. He is currently working towards the PhD degree in Electrical and Computer Engineering Department, Colorado State University, Fort Collins,

CO, USA. His research interests include embedded system, systems-on-chip and optical networks-on-chip. He is a Student Member of the IEEE.

Sudeep Pasricha is an Associate Professor in the Electrical and Computer Engineering Department, Colorado State University, Fort Collins, CO, USA. His research interests include embedded, mobile, and high-performance computing. He received a PhD degree in computer science from the University of California Irvine, CA, USA in 2008. He is a Senior Member of the IEEE and ACM.

■ Direct questions and comments about this article to Sai Vineel Reddy Chittamuru, Electrical Computer Engineering, Colorado State University, Fort Collins, CO, 80523-1373; sai.chittamuru@colostate.edu.