

# Process Variation Aware Crosstalk Mitigation for DWDM based Photonic NoC Architectures

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

## Abstract

Photonic network-on-chip (PNoC) architectures are a potential candidate for communication in future chip multiprocessors as they can attain higher bandwidth with lower power dissipation than electrical NoCs. PNoCs typically employ dense wavelength division multiplexing (DWDM) for high bandwidth transfers. Unfortunately, DWDM increases crosstalk noise and decreases optical signal to noise ratio (SNR) in microring resonators (MRs) threatening the reliability of data communication. Additionally, process variations induce variations in the width and thickness of MRs causing shifts in resonance wavelengths of MRs, which further reduces signal integrity, leading to communication errors and bandwidth loss. In this paper, we propose a novel encoding mechanism that intelligently adapts to on-chip process variations, and improves worst-case SNR by reducing crosstalk noise in MRs used within DWDM-based PNoCs. Experimental results on the Corona PNoC architecture indicate that our approach improves worst-case SNR by up to 44.13%.

## 1. Introduction

The ever increasing demand for higher performance computing and aggressive technology scaling have driven the trend of integrating a steadily increasing number of processing cores on a single die. With increase in core count, electrical networks-on-chip (ENoCs) are projected to suffer from crippling high power dissipation and severely reduced performance [1]-[3]. Recent developments in the area of silicon photonics have enabled the integration of on-chip photonic interconnects with CMOS circuits, and led to the introduction of photonic networks-on-chip (PNoCs) that can offer ultra-high bandwidth, reduced power dissipation, and lower latency than ENoCs.

Several crossbar topology based PNoC architectures have been proposed to date (e.g., [4]-[6]). These architectures are built using silicon photonic devices such as microring-resonators (MRs) and silicon waveguides, and employ dense-wavelength-division-multiplexing (DWDM), where a large number of wavelengths are multiplexed in a waveguide to enable high bandwidth parallel data transfers. Unfortunately, the deleterious effects of fabrication process variations (PV) in MRs and silicon waveguides can reduce reliability in PNoCs. The PV-induced variations in the width and thickness of MRs cause resonance wavelength shifts in MRs [7]-[8]. An MR couples light of a specific resonance wavelength to/from the waveguide, enabling electrical-to-photonic (modulation) and photonic-to-electrical (detection) operations. PV-induced resonance shifts worsen the signal-to-noise-ratio (SNR) in MRs, as they decrease the signal power and increase the crosstalk noise power. This in turn deteriorates the bit-error-rate (BER) in a waveguide. For example, a previous study shows that in a DWDM-based photonic interconnect, when PV-induced resonance shift is over 1/3 of the channel gap, BER of photonic transmission increases from  $10^{-12}$  to  $10^{-6}$  [10]. The increase in BER decreases the reliability of data communication.

The most intuitive way to counteract PV-induced resonance shifts in MRs is to realign the resonant wavelengths by using localized trimming [9] and thermal tuning [7] mechanisms. The former approach causes the wavelength to shift towards the blue end and the latter towards the red end of the resonance spectrum. These mechanisms improve SNR and reduce BER in DWDM based photonic interconnects. However, our analysis in this paper shows that the use of localized trimming methods increases the intrinsic optical loss in the MRs [23][25] and waveguide due to the free carrier absorption effect (FCA). This additional signal loss decreases Q-factor of MRs, which in turn increases crosstalk noise in detectors and reduces SNR. Thus, *the use of trimming to remedy PV is not a viable option, strongly motivating new crosstalk mitigation techniques in PV-affected PNoCs.*

We observe that while transmitting data in DWDM-based PNoCs, the impact of PV remedial techniques (such as localized trimming) on the crosstalk noise in MRs depends not only on the amount of PV-induced resonance shifts but also on the characteristics of data values propagating in the network. This implies that the harmful effects of localized trimming on crosstalk noise can be reduced by controlling the relative occurrences of some data values. Therefore, in this paper, we present a novel technique that intelligently reduces undesirable data value occurrences in a photonic waveguide based on the process variation profile of MRs in the detecting nodes. Our technique is easily implementable in any existing DWDM-based crossbar PNoC without requiring major modifications to the architectures. Further, our technique is lightweight and possesses low overhead. To the best of our knowledge, this is the first work that attempts to improve SNR in PNoCs considering PV in its MRs. Our novel contributions in this work are summarized below:

- We present a device-level analytical model that captures the deleterious effects of localized trimming in MRs. Moreover, we extend this model for system level crosstalk analysis;
- We propose a double bit crosstalk mitigation mechanism (DBCTM) that improves the worst-case SNR of the PV affected detecting node by encoding specific portions of data to avoid undesirable data occurrences;
- We evaluate our proposed technique by implementing it on the well-known crossbar PNoC Corona architecture [4], [11], and compare it with two previously proposed encoding based crosstalk mitigation mechanisms from [13] for real-world multi-threaded PARSEC [12] benchmarks.

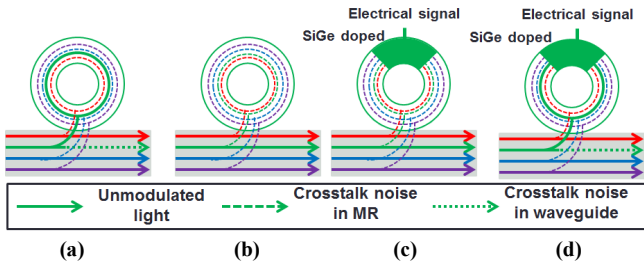
## 2. Related Work

DWDM-based PNoCs utilize photonic devices such as microring resonators (MRs) as modulators and detectors, photonic waveguides, splitters, and trans-impedance amplifiers (TIAs). The reader is directed to [13] for more information on these devices. Each constituent photonic device in a PNoC contributes to some type of optical signal loss, the combined effect of which

negatively affects the system SNR. In addition to the optical signal loss, the crosstalk noise of the constituent MRs also deteriorates the SNR. Crosstalk is an intrinsic property of every MR, so both modulators and detectors are susceptible to crosstalk noise in DWDM-based PNoCs. Fig. 1 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 wavelength by removing the light pulse, there is crosstalk generated in the waveguide, as shown in Fig. 1 (a) and 1(d).

One of the key challenges for the widespread adoption of DWDM-based PNoC architectures is to mitigate the crosstalk noise in their MRs. The effect of crosstalk noise on SNR is negligible in DWDM systems presented in [14] and [15], as these systems use only four DWDM wavelengths per waveguide. On the other hand, in crossbar architectures such as Corona [4] that use 64 wavelength DWDM, there exists significant crosstalk noise. The quantitative results in [16] demonstrate the damaging impact of crosstalk noise in Corona, where the worst-case SNR is estimated to be about 14dB in data waveguides, which is insufficient for reliable data communication. To mitigate the impacts of crosstalk noise in DWDM based PNoCs, two encoding techniques PCTM5B and PCTM6B are presented in [13]. In [9] a technique was proposed to increase channel spacing between adjacent wavelengths in DWDM, to mitigate crosstalk in MR detectors. However, *none of these works considers process variations and their impact on crosstalk in DWDM-based PNoCs.*

A few prior works have explored the impact of process variations on DWDM-based photonic links [10][25]. In [10], the authors discuss a run-time hardware-software management solution that optimizes the performance and reliability of photonic communication to compensate for PV effects. In [25], a methodology to salvage network-bandwidth loss due to process-variation-drifts is proposed, which reorders MRs and trims them to nearby wavelengths. *All of these PV-remedial techniques are network specific and ignore the harmful effects of PV on crosstalk.* In contrast, we consider the deleterious effects of PV-remedial techniques and propose a generalized technique for crosstalk noise mitigation with minimal overhead, to improve SNR and communication reliability in DWDM-based crossbar PNoCs.



**Fig. 1: MR operation phases in DWDM-based waveguides (note: MR shown has green resonant wavelength): (a) modulator modulating in resonance wavelength (b) modulator in passing (through) mode (c) detector in passing mode (d) detector in detecting mode.**

### 3. Analytical Model for PV-Aware Crosstalk Analysis

#### 3.1 Impact of localized trimming on crosstalk

As discussed earlier, the localized trimming method is essential to deal with PV-induced resonance red shifts in MRs. However, the use of this method in an MR alters its intrinsic optical properties, which leads to increased crosstalk noise and degraded performance in PNoCs using these MRs. In this section, we discuss the effects of the localized trimming method and present analytical models to capture these effects in MRs. Further, we

extend these models to generate system-level models for the Corona PNoC architecture [4]. The system-level models enable quantification of signal and noise powers in the constituent MRs and DWDM waveguides of the Corona architecture.

An MR can be considered to be a circular photonic waveguide with a small diameter (Fig. 1), not to be confused with the larger DWDM-based photonic waveguide for which MRs serve as modulators and detectors. The localized trimming method injects extra free carriers in the circular MR waveguide to counteract the PV-induced resonance red shifts in it. The introduction of extra free carriers reduces the refractive index of the circular MR waveguide, which in turn induces a blue shift in resonance to counteract the PV-induced red shifts. However, the extra free carriers increase the absorption related optical loss in the MR due to the free carrier absorption effect (FCA) [23]. The increase in the intrinsic optical loss results in a decrease of MR Q-factor, which in turn increases the MR insertion loss and crosstalk.

We use a process variation map (discussed in Section 3.3) to estimate PV-induced shifts in the resonance wavelengths of all the MRs across a chip. Then, for each MR device, we calculate the amount of change in refractive index ( $\Delta n_{si}$ ) required to counteract this wavelength shift using the following equation [22]:

$$\Delta n_{si} = \frac{\Delta \lambda_r \cdot n_g}{\Gamma \cdot \lambda_r}, \quad (1)$$

where,  $\Delta \lambda_r$  is the PV-induced resonance shift that need to be compensated for,  $\lambda_r$  is the target resonance wavelength,  $n_g$  is the group refractive index of the MR waveguide, and  $\Gamma$  is the confinement factor describing the overlap of the optical mode with the MR waveguide’s silicon core. We assume that the MR waveguides used in this study are similar to those reported in [23], fabricated using standard Si-SiO<sub>2</sub> material with a cross section of 450nm×250nm. The values of  $\Gamma$  and  $n_g$ , for these MR waveguides, are set to 0.7 and 4.2 respectively [23].

The required change in the free carrier concentration to induce the refractive index change of  $\Delta n_{si}$  at around 1.55μm wavelength can be quantified using the following equation [22]:

$$\Delta n_{si} = -8.8 \times 10^{-22} \Delta N_e - 8.5 \times 10^{-18} (\Delta N_h)^{0.8}, \quad (2)$$

where,  $\Delta N_e$  and  $\Delta N_h$  are the change in free electron concentration and the change in free hole concentration respectively. The change in the absorption loss coefficient ( $\Delta \alpha_{si}$ ) due to the change in free carrier concentration (owing to the FCA effect) can be quantified using the following equation [22]:

$$\Delta \alpha_{si} = -8.5 \times 10^{-18} \Delta N_e - 6.0 \times 10^{-18} \Delta N_h, \quad (3)$$

The Q-factor of an MR depends on the absorption loss coefficient. The relation between the Q-factor and  $\Delta \alpha_{si}$ , assuming critical coupling of MRs, is given by the following equation [23], where Q’ is the loaded q-factor of MR.

$$Q' = Q + \Delta Q = \frac{\pi n_g}{\lambda_r (\alpha + \Delta \alpha_{si})}, \quad (4)$$

where,  $\Delta Q$  is the change in Q-factor and  $\alpha$  is the original loss coefficient, which is a sum of three components: (i) intrinsic loss coefficient due to material loss and surface roughness; (ii) bending loss coefficient, which is a result of the curvature in the MR; and (iii) the absorption effect factor that depends on the original free carrier concentration in the waveguide core. Typically, the localized trimming method injects excess concentration of free carriers into the MR, which leads to an increase in the absorption loss coefficient (positive  $\Delta \alpha_{si}$ ). As evident from Eq. (4), a positive value of  $\Delta \alpha_{si}$  results in decrease of Q-factor. This causes a broadening of the MR passband, which results in increased insertion loss and crosstalk power penalties.

### 3.2 Crosstalk modeling for Corona PNoC

In this work, we characterize crosstalk in DWDM-waveguides for the well-known Corona PNoC enhanced with token-slot arbitration [4], [11]. In DWDM-based waveguides, data transmission requires modulating light using a group of MR modulators equal to the number of wavelengths supported by DWDM. Similarly, data detection at the receiver requires a group of detector MRs equal to the number of DWDM wavelengths. We present analytical equations that model worst-case crosstalk noise power, maximum power loss, and SNR in detector MR groups (similar equations are applicable to modulator MR groups).

**Table I: Notations for photonic power loss, crosstalk coefficients [16]**

Notation	Parameter type	Parameter value (in dB)
$L_p$	Propagation loss	-0.274 per cm
$L_B$	Bending loss	-0.005 per 90°
$L_{MI}$	Inactive modulator through loss	-0.005
$L_{MA}$	Active modulator power loss	-0.6
$L_{DP}$	Passing detector through loss	-0.005
$L_{DD}$	Detecting detector power loss	-1.6
$L_{S12}$	1X2 splitter power loss	-0.2
$L_{S14}$	1X4 splitter power loss	-0.2
$L_{S15}$	1X5 splitter power loss	-0.2
$L_{S16}$	1X6 splitter power loss	-0.2
$X_{MA}$	Active modulator crosstalk factor	-16
$X_{DD}$	Detecting detector crosstalk factor	-16

**Table II: Other model parameter notations**

Notation	Crosstalk Coefficient	Parameter Value
$Q$	Q-factor	9000
FSR	Free spectral range	51.2nm
$L$	Photonic path length in cm	
$B$	Number of bends in photonic path	
$\lambda_j$	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	

Before presenting actual analytical equations, we provide notations for different parameters used in the analytical equations in Tables I and II. The Corona PNoC is designed for a 256 core single-chip platform, where cores are grouped into 64 clusters, with 4 cores in each cluster. A photonic crossbar topology with 64 data channels is used for communication between clusters. Each channel consists of 4 multiple-write-single-read (MWSR) waveguides with 64-wavelength 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.

A data channel starts at a cluster called ‘home-cluster’, traverses other clusters (where modulators can modulate light 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 from an off-chip laser to each of the 64 data channels at its home-cluster, through a series of 1X2 splitters. In each of the 64 home-clusters, optical power is distributed among 4 MWSR waveguides equally using a 1X4 splitter with splitting factor  $R_{S14}$ . As all 1X2 splitters are present before the last (64<sup>th</sup>) channel, this channel suffers the highest signal power loss. Thus, the worst-case signal and crosstalk noise exists in the detector group of the 64<sup>th</sup> cluster node, and this node is defined as the worst-case power loss node ( $N_{WCPL}$ ) in Corona. For this  $N_{WCPL}$  node, the signal power ( $P_{signal}(j)$ ) and crosstalk noise power ( $P_{noise}(j)$ ) received at each detector  $j$  are expressed in Eq. (5) and (6) [14].  $P_S(i,j)$  in Eq. (7) is the signal power of the  $i^{th}$  wavelength received before the  $j^{th}$  detector. Similarly in Eq.

(9),  $P_N(i,j)$  is the crosstalk noise power of the  $i^{th}$  wavelength before the  $j^{th}$  detector.  $K_S$  and  $K_N$  in Eq. (10) and (11) represent signal and crosstalk noise power losses before the detector group of  $N_{WCPL}$ ;  $\psi(i,j)$  in Eq. (8a) represents signal power loss of the  $i^{th}$  wavelength before the  $j^{th}$  detector within the detector group of  $N_{WCPL}$ ; and  $\Phi(i,j)$  in Eq. (8b) is the crosstalk coupling factor of the  $i^{th}$  wavelength and the  $j^{th}$  detector. Due to the use of trimming to remedy PV, crosstalk coupling factor ( $\Phi$ ) increases with decrease in loaded Q-factor ( $Q'$ , Eq. (4)), which in turn increases crosstalk noise in the detectors. For instance, realignment of PV-induced resonance shift of 0.05nm ( $\Delta\lambda_r$ , Eq. (1)) using the trimming method increases the absorption loss coefficient ( $\alpha + \Delta\alpha_{Si}$ ) to 11.5cm<sup>-1</sup> from 9.5cm<sup>-1</sup> ( $\alpha$ , corresponding to the original Q-factor of 9000 given in Table II). This increase in  $\alpha$  reduces the Q-factor to 7400 from 9000 (Eq. (4)), which in turn increases  $\delta$  (Eq. (8b)) by 21.6% resulting in about 15 $\times$  increase in the coupling factor  $\Phi$  (Eq. (8b)). As evident from Eq. (6), this increase in  $\Phi$  results in the increase of crosstalk noise in MR detectors.

We can define SNR( $j$ ) of the  $j^{th}$  detector of  $N_{WCPL}$  as the ratio of  $P_{signal}(j)$  to  $P_{noise}(j)$ , as shown in Eq. (12). 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 the data waveguide. Section 4 uses these models to explain how crosstalk mitigation techniques impact SNR.

$$P_{signal}(j) = L_{DD} P_S(j, j) \quad (5)$$

$$P_{noise}(j) = L_{DD} P_N(j, j) + \sum_{i=1}^n \Phi(i, j) (P_S(i, j) + P_N(i, j)) \quad (i \neq j) \quad (6)$$

$$P_S(i, j) = K_S \psi(i, j) P_{in}(i) \quad (7)$$

$$\psi(i, j) = \begin{cases} X_{DD}(L_{DP})^{j-1}, & \text{If } j - 1 \geq i \text{ and } D_B = 1 \\ (L_{DP})^{j-1}, & \text{if } j - 1 < i \text{ and } D_B = 1 \\ X_{MA}X_{DD}(L_{DP})^{j-1}, & \text{If } j - 1 \geq i \text{ and } D_B = 0 \\ X_{MA}(L_{DP})^{j-1}, & \text{If } j - 1 < i \text{ and } D_B = 0 \end{cases} \quad (8a)$$

$$\Phi(i, j) = \frac{\delta^2}{((i-j)\frac{FSR}{n})^2 + \delta^2}, \text{ Here } \delta = \frac{\lambda_j}{2Q'} \quad (8b)$$

$$P_N(i, j) = \begin{cases} 0, & \text{If } j > i \text{ and } D_B = 1 \\ K_N(L_{DP})^j P_{in}(i), & \text{if } j \leq i \text{ and } D_B = 1 \\ 0, & \text{If } j > i \text{ and } D_B = 0 \\ X_{MA}K_N(L_{DP})^j P_{in}(i), & \text{if } j \leq i \text{ and } D_B = 0 \end{cases} \quad (9)$$

$$K_S = \begin{cases} (R_{S14})(L_{S14})(L_p)^L (L_B)^B (L_{MI})^{64 \times 63}, & \text{If } D_B = 1 \\ (R_{S14})(L_{S14})(L_p)^L (L_B)^B (L_{MI})^{64 \times 62 + 63}, & \text{If } D_B = 0 \end{cases} \quad (10)$$

$$K_N = \begin{cases} (R_{S14})(L_{S14})(L_p)^L (L_B)^B X_{MA} (L_{MI})^{64 \times 62 + 63}, & \text{If } D_B = 1 \\ (R_{S14})(L_{S14})(L_p)^L (L_B)^B X_{MA} (L_{MI})^{64 \times 62 + 62}, & \text{If } D_B = 0 \end{cases} \quad (11)$$

$$SNR(j) = \frac{P_{signal}(j)}{P_{noise}(j)} \quad (12)$$

### 3.3. Modeling PV of MR devices in Corona PNoC

We adapt the VARIUS tool [24] to model die-to-die (D2D) as well as within-die (WID) process variations in MRs for the Corona PNoC. We consider photonic devices with a silicon (Si) core and silicon-dioxide (SiO<sub>2</sub>) cladding. VARIUS uses a normal distribution to characterize on-chip D2D and WID process variations. The key parameters are mean ( $\mu$ ), variance ( $\sigma^2$ ), and density ( $\alpha$ ) of a variable that follows the normal distribution. As wavelength variations are approximately linear to dimension variations of MRs, we assume they follow the same distribution. The mean ( $\mu$ ) of wavelength variation of an MR is its nominal resonance wavelength. We consider a DWDM wavelength range in the C and L bands [13], with a starting wavelength of 1550nm and a channel spacing of 0.8nm. Hence, those wavelengths are

the means for each MR modeled. The variance ( $\sigma^2$ ) of wavelength variation is determined based on laboratory fabrication data [7] and our target die size. We consider a 256-core chip with die size 400 mm<sup>2</sup> at 22nm CMOS technology node. For this die size we consider a WID standard deviation ( $\sigma_{WID}$ ) as 0.61nm [25] and D2D standard deviation ( $\sigma_{D2D}$ ) as 1.01 nm [25]. Further we also consider a density ( $\alpha$ ) of 0.5 [25] for this die size. With these parameters, we use VARIUS to generate 100 process variation maps. Each process variation map contains over one million points indicating the PV-induced resonance shift of MRs. The total number of points picked from these maps are equal to the number of MRs in the Corona PNoC architecture.

## 4. Double-bit Crosstalk Mitigation (DBCTM)

### 4.1. Overview

In this section, we present our proposed double-bit crosstalk mitigation (DBCTM) technique, which is illustrated in Figure 2. As evident from Eq. (6), crosstalk noise in MR detectors of DWDM-based PNoCs increases with increase in coupling factor ( $\Phi$ ) and increase in signal strength of immediate non-resonating wavelength. This implies that the crosstalk noise in a detector can be reduced by reducing the signal strength of immediate non-resonating wavelengths. Therefore, our proposed DBCTM technique decreases the signal strength of immediate non-resonant wavelength by modulating a zero on it, which results in reduced crosstalk noise in the detector. For that, the DBCTM technique first divides detecting MRs into groups of 8 MRs each. Then, it determines the maximum PV-induced resonance red shift ( $\Delta\lambda_{max}$ ) in each MR group. The PV-induced resonance shifts in MRs can be gauged by measuring the variations in the thickness and width values of MRs. This is because there is a linear dependency between thickness and width variations of MRs and the resonance wavelength drift of MRs, as discussed in [7] and [10]. For example, a 1nm variation in width and height of an MR leads to 0.58~1nm [7], [10] and ~2nm [7] resonance shifts respectively. In a real system, spectroscopic ellipsometry (SE) and critical dimension scanning electron microscopy (CD-SEM) can be employed to measure thickness and width variations in actual fabricated MR devices in an automated manner [7]. In our analysis, we model and estimate PV in MRs using the VARIUS tool [24], a description of which is given in Section 3.3.




Fig. 2: Overview of proposed DBCTM technique

Once PV-induced resonance red shifts of MRs are determined, we store information about whether to enable or disable encoding for each MR group in a read-only memory (ROM) at the modulating node, based on the maximum PV-induced resonance red shift ( $\Delta\lambda_{max}$ ) value for the group. If this value is greater than a threshold red shift value ( $\Delta\lambda_{th}$ ) for an MR group, we store a ‘1’ to enable DBCTM, else we store a ‘0’ to disable DBCTM for this MR group. MR groups with  $\Delta\lambda_{max} < \Delta\lambda_{th}$  are thus not impacted. Only MR-groups with  $\Delta\lambda_{max} > \Delta\lambda_{th}$  employ encoding.

### 4.2 DBCTM sensitivity analysis with Corona PNoC

Our encoding scheme involves injecting zeros between data

bits. These extra bits are called shielding bits. As the number of shielding bits increases, laser power and trimming power of PNoCs also increase. Thus, we need to limit the number of shielding bits. We performed a sensitivity analysis using the Corona architecture with varying number of shielding bits per detecting node to quantify its effect on worst-case SNR. We analyzed worst case SNR with 0%, 25%, 50%, 75% and 100% of shielding bits added to data bits for the Corona PNoC. Based on our analysis across 100 process variation maps (see section 3.3), we determined  $\Delta\lambda_{th}$  to be 0.45nm, 0.88 nm, 1.25nm and 4.25nm, for the cases with 25%, 50%, 75% and 100% of shielding bits to data bits, respectively. Fig. 3 shows the range of worst-case SNR values across process variation maps, for different ratios of shielding bits to data bits. From the figure it can be seen that on an average DBCTM with 25%, 50%, 75% and 100% shielding bits has 8.1%, 19.67%, 26% and 40.5% higher worst case SNR (note: higher SNR is better) respectively compared to the baseline (with 0% shielding). Intuitively, higher ratios of shielding bits to data bits should result in higher worst case SNR, as more shielding bits can be used to shield data bits, which in turn reduces crosstalk noise and improves SNR. But, with increase in number of shielding bits, the number of MRs on the waveguides increases. High MR count on the waveguide results in higher through losses, which increases laser power. Compensating for PV drifts of high MR counts requires high trimming/tuning power in PNoCs. Fig. 3 shows that average power consumption of DBCTM with 25%, 50%, 75% and 100% shielding bits is 14%, 20.1%, 63.9% and 104.1% higher compared to the baseline. To balance crosstalk reliability and power overheads, we select the 50% shielding bits to data bits configuration, for the rest of our experiments.




Fig. 3: Sensitivity analysis in terms of worst-case SNR of Corona architecture with DBCTM allowing 0%, 25%, 50% and 100% ratio of shielding bits to data bits across 100 process variation maps; average power consumption are also shown on the top of each bar.

To implement our DBCTM technique with 50% shielding bits on the Corona PNoC, we increase the number of MWSR waveguides in each channel from 4 to 6, to maintain the same bandwidth as in the baseline case. To distribute optical power between these waveguides, there is also a need to replace 1X4 splitters with 1X6 splitters with a splitting factor of  $R_{S16}$ . Additionally each modulating node needs to store 2,646 bits in its ROM to capture encoding requirements for all the remaining 63 detecting nodes. Lastly, we consider up to a two cycle overhead for encoding and decoding of data in DBCTM, as per our implementations at 5GHz. The first cycle is needed to retrieve data from ROM, whereas the second cycle is used if data is to be encoded before sending on the waveguide.

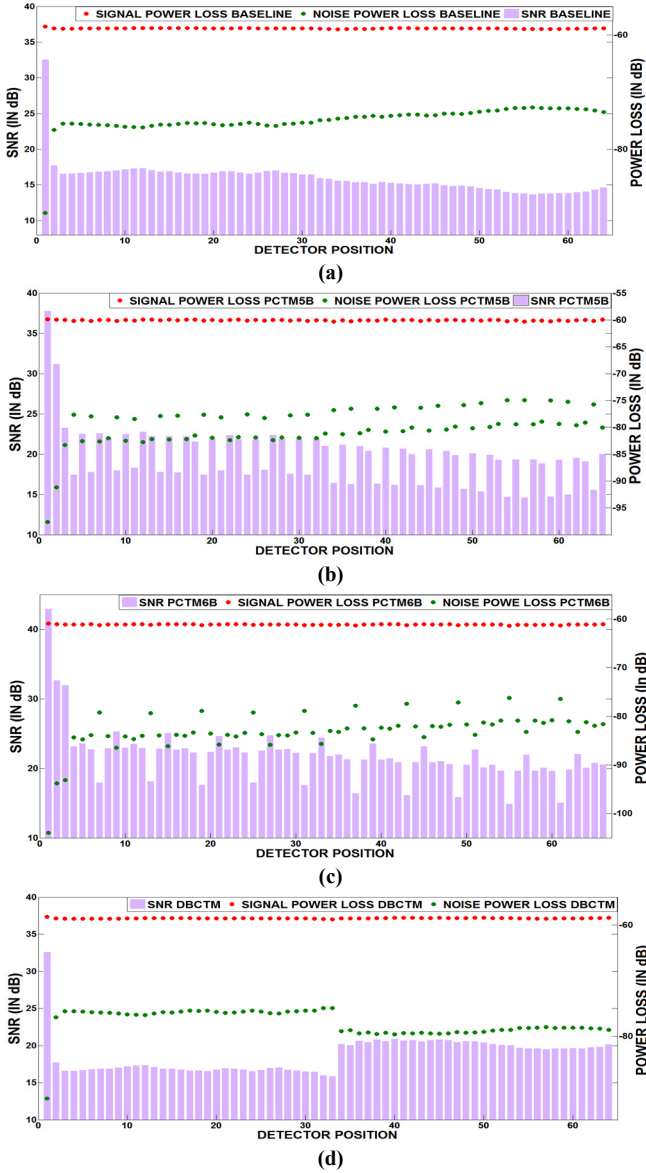
## 5. Experiments

### 5.1. Experimental setup

To evaluate our proposed crosstalk noise mitigation technique (DBCTM) in DWDM-based PNoCs, we implement and integrate it for the Corona [4], [16] crossbar-based PNoC architecture. We



modeled and performed simulation based analysis of the enhanced Corona PNoC using a cycle-accurate NoC simulator, for a 256 core single-chip architecture at 22nm. As explained in the previous section we generated 100 process variation maps to evaluate the impact of PV on our DBCTM technique. We used real-world traffic from applications in the PARSEC benchmark suite [12]. GEM5 full-system simulation [17] of parallelized PARSEC applications was used 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 20mm×20mm chip size, to determine lengths of MWSR waveguides in the Corona architecture. Based on this analysis, we estimated the time needed for light to travel from the first to the last node as 8 cycles at 5 GHz clock frequency. We use a 512 bit packet size, as advocated in the Corona PNoC.



**Fig. 4: Detector-wise signal power loss, crosstalk noise power loss and minimum SNR in worst-case power loss node for one process variation map of Corona (a) baseline with 64-detectors (b) PCTM5B with 65-detectors (c) PCTM6B with 66-detectors (d) DBCTM with 64-detectors.**


The static and dynamic energy consumption of electrical routers and concentrators in Corona is based on results from the

Orion 2.0 [18] tool. We model and consider area, power, and performance overheads for DBCTM (50% case). The electrical area overhead is estimated to be 6.24 mm<sup>2</sup> and power overhead is estimated to be 1.14 W, using gate-level analysis and the CACTI 6.5 [19] tool for memory and buffers. The photonic area overhead is estimated to be 9.12 mm<sup>2</sup> based on the physical dimensions [15] of waveguides, MRs, and splitters. For energy consumption of photonic devices, we adopt model parameters from recent work [16], [20], [21], with 0.42pJ/bit for every modulation and detection event and 0.18pJ/bit for the driver circuits of modulators and photodetectors. We used optical loss for photonic components, as shown in Table I, to determine the photonic laser power budget and correspondingly the electrical laser power.

## 5.2. Experimental results with Corona PNoC

Utilizing the models presented in section 3, we calculate the received crosstalk noise and SNR at detectors for the node with worst-case power loss ( $N_{WCPL}$ ), which corresponds to MR detectors in cluster 64 for the Corona architecture.

Our first set of experiments compares the baseline Corona architecture with fair token-slot arbitration [4], [16] but without any crosstalk-enhancements, with three variants of the architecture corresponding to the three crosstalk-mitigation strategies we compare: PCTM5B and PCTM6B proposed in [13] and our proposed DBCTM technique from this paper. The worst-case SNR 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 PCTM5B, PCTM6B and DBCTM crosstalk-mitigation techniques, this is not the case, i.e., each detector in cluster 64 has a worst-case SNR 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, for an accurate analysis.



**Fig. 5: Worst-case SNR comparison of DBCTM with PCTM5B [13] and PCTM6B [13] for Corona PNoC considering 100 process variation map**

Fig. 4 (a)-(d) present detector signal power loss, crosstalk noise power loss, and SNR corresponding to the detectors in the 64<sup>th</sup> cluster for the baseline and three variants of the Corona architecture for a randomly selected PV map generated using VARIUS, as explained in section 3.3. Fig. 4(b) indicates that worst-case SNR (lowest value of the bars, which represent SNR in detectors) improves notably over the baseline in Fig. 4(a) when using the PCTM5B technique. However, the improvement is on the lower side for the remaining detectors. Fig. 4(c) shows that the PCTM6B technique improves worst-case SNR marginally over PCTM5B, but does a better job of improving SNR significantly for most of the detectors. Fig. 4(d) shows that the DBCTM technique results in a significant improvement in worst-case SNR as well as SNR for all detectors compared to the baseline, as well as the PCTM5B and PCTM6B techniques. The worst-case SNR is obtained at the 42<sup>nd</sup> detector of the 64<sup>th</sup> cluster in the baseline case; whereas for the PCTM5B, PCTM6B and DBCTM configurations worst-case SNR occurs at the 59<sup>th</sup>, 61<sup>st</sup> and 33<sup>rd</sup> detectors of the same cluster, respectively.

Fig. 5 summarizes the worst-case SNR results for the baseline, PCTM5B, PCTM6B, and DBCTM techniques. From the figure, it can be surmised that Corona with DBCTM has 19.28-44.13%, 12.44-34.19% and 4.5-31.30% worst-case SNR improvements on average, compared to baseline, PCTM5B, and PCTM6B respectively. Both the PCTM5B and PCTM6B techniques eliminate occurrences of ‘111’ in a data word and have limited occurrences of ‘11’, which helps reduce crosstalk noise in the detectors. But, these techniques do not consider the impact of PV resonance wavelength drifts, which leads to worse SNR degradation. The DBCTM technique reduces crosstalk noise in the detectors by using shielding bits between data bits, further it also considers the PV profile of MRs to select MRs for shielding.

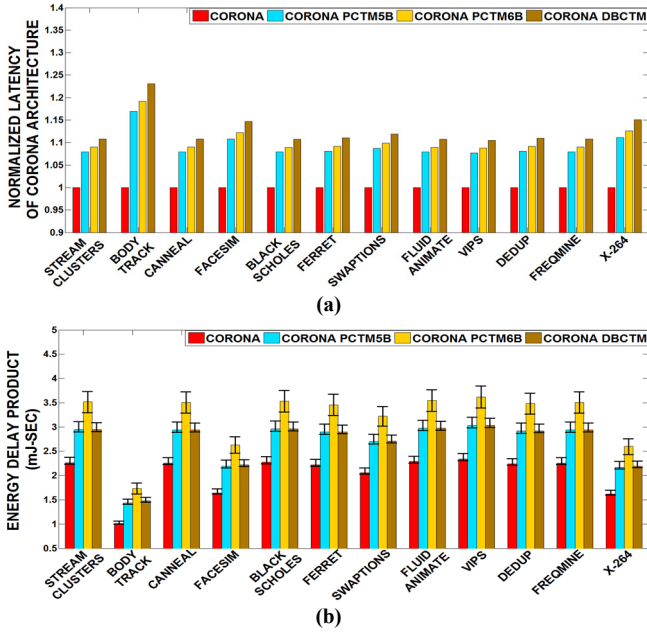


Fig. 6: (a) normalized latency (b) energy-delay product (EDP) comparison between Corona baseline and Corona architecture with PCTM5B, PCTM6B and DBCTM techniques, with PARSEC benchmarks. All results are normalized to the baseline Corona architecture results.

Fig. 6 (a) and (b) present simulation results that quantify the average network packet latency and energy-delay product (EDP) for the four Corona configurations. Results are shown for twelve multi-threaded PARSEC benchmarks. From Fig. 6(a) it can be seen that on average, the Corona configuration with DBCTM has 12.6%, 3.4% and 2.1% higher latency compared to baseline, PCTM5B and PCTM6B respectively. The additional delay due to encoding and decoding of data with DBCTM, PCTM5B and PCTM6B contributes to their increase in average latency. The penalty due to encoding/decoding is 1 cycle in PCTM5B and PCTM6B, whereas DBCTM has a 2 cycle penalty which increases its delay overhead. From the results for EDP shown in Fig. 6(b), it can be seen that on average, Corona with the DBCTM technique has 31.6% higher EDP compared to the baseline. 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. This in turn increases static energy consumption. Dynamic energy also increases in these architectures, but by much less. However, EDP for the DBCTM technique is 16.4% lower compared to PCTM6B. Despite the higher latency overhead compared to PCTM6B, DBCTM saves considerable laser and trimming/tuning power due to lower photonic hardware requirements than PCTM6B.

## 6. Conclusion

We have presented a crosstalk mitigation technique for the reduction of crosstalk noise in the detectors of dense wavelength division multiplexing (DWDM) based photonic network-on-chip (PNoC) architectures with crossbar topologies. Our DBCTM technique show interesting trade-offs between reliability, performance, and energy overhead for the Corona crossbar PNoC architectures. Our experimental analysis shows that the DBCTM technique improve worst-case SNR by 44.13% compared to the baseline Corona architecture, and by 31.30% compared to the best known PNoC crosstalk mitigation scheme from prior work.

## 7. Acknowledgments

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

## References

- [1] L. Benini and G. De Micheli, “Networks on chip: A new paradigm for systems on chip design,” in Proc. DATE, 2002, pp. 418–419.
- [2] J. D. Owens, et al., “Research challenges for on-chip interconnection networks,” in IEEE Micro, September-October 2007.
- [3] (2013). International Technology Roadmap for Semiconductors (ITRS) [Online]. Available: <http://www.itrs.net>
- [4] D. Vantrease et al., “Corona: System implications of emerging nanophotonic technology,” in Proc. ISCA, 2008.
- [5] Y. Pan et al., “Firefly: Illuminating future network-on-chip with nanophotonics,” in Proc. ISCA, 2009.
- [6] S. V. R. Chittamuru et al., “A Reconfigurable Silicon-Photonic Network with Improved Channel Sharing for Multicore Architectures,” in ACM GLSVLSI, May, 2015.
- [7] S. K. Selvaraja., “Wafer-Scale Fabrication Technology for Silicon Photonic Integrated Circuits,” PhD thesis, Ghent University, 2011.
- [8] J. Orcutt et al., “Nanophotonic integration in state-of-the-art cmos foundries,” Optics Express, 19:2335–2346, 2011.
- [9] S. V. R. Chittamuru et al., “Improving Crosstalk Resilience with Wavelength Spacing in Photonic Crossbar-based Network-on-Chip Architectures,” IEEE MWSCAS, Aug. 2015.
- [10] Z. Li et al., “Reliability modeling and management of nanophotonic on-chip networks,” IEEE TVLSI, 20:98–111, 2010.
- [11] D. Vantrease et al., “Light speed arbitration and flow control for nanophotonic interconnects,” in Proc. IEEE/ACM MICRO, 2009.
- [12] C. Bienia et al., “The PARSEC Benchmark Suite: Characterization and Architectural Implications,” in PACT, Oct. 2008.
- [13] S. V. R. Chittamuru et al., “Crosstalk Mitigation for High-Radix and Low-Diameter Photonic NoC Architectures,” in IEEE D&T, 2015.
- [14] Q. Xu, B. Schmidt et al., “Cascaded silicon micro-ring modulators for wdm optical interconnection,” in Opt. Exp., vol. 14, 2006.
- [15] S. Xiao et al., “Modeling and measurement of losses in silicon-on-insulator resonators and bends,” in Opt. Exp., vol.15, no.17, 2007.
- [16] L.H.K. Duong et al., “A Case Study of Signal-to-Noise Ratio in Ring-Based Optical Networks-on-Chip,” in IEEE D&T, 2014.
- [17] N. Binkert et al., “The gem5 Simulator,” in CA News, May 2011.
- [18] A.B. Kahng et al., “ORION 2.0: A Power-Area Simulator for Interconnection Networks,” in IEEE Trans. on VLSI 20, 2012, 191–196.
- [19] CACTI 6.5, <http://www.hpl.hp.com/research/cacti/>
- [20] X. Zheng et al., “Ultra-efficient 10Gb/s hybrid integrated silicon photonic transmitter and receiver,” in Opt. Express, Mar 2011.
- [21] P. Grani and S. Bartolini, “Design Options for Optical Ring Interconnect in Future Client Devices,” in ACM JETC, May, 2014.
- [22] R. G. Beausoleil, “Large-Scale Integrated Photonics for High-Performance Interconnects,” ACM JETC, Vol. 7, No. 2, 2011.
- [23] K. Preston, et al., “Performance guidelines for WDM interconnects based on silicon microring resonators,” IEEE CLEO, 2011.
- [24] S. Sarangi et al., “Varius: A model of process variation and resulting timing errors for microarchitects,” IEEE TSM, 21(1):3–13, 2008.
- [25] Y. Xu et al., “Tolerating process variations in nanophotonic on-chip networks,” in Proc. ISCA, Portland, OR, USA, 2012, pp. 142–152.
- [26] J. Karttunen et al., “Loading effects in deep silicon etching,” in International Society of Optical Engineering, vol. 4174, pp. 90–97, 2000.