

Mitigation of Homodyne Crosstalk Noise in Silicon Photonic NoC Architectures with Tunable Decoupling

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 signal modulation, to enable high bandwidth on-chip transfers. Unfortunately, due to the resonant nature of MRs, the power built-up in their cavity gradually recouples back into the photonic waveguides. This recoupled power induces time-dependent unfilterable homodyne crosstalk noise, when the wavelength of the recoupled power matches with the wavelength of a signal in the waveguide. The homodyne crosstalk in turn deteriorates the signal-to-noise ratio (SNR) and on-chip communication reliability. This paper presents a novel lightweight technique to mitigate homodyne crosstalk noise in DWDM-based PNoCs. We evaluate the effectiveness and overhead of our technique by implementing it for well-known PNoC architectures, including Corona, Firefly and Flexishare. Experimental results indicate that our approach when implemented on these PNoCs can improve the worst-case SNR by up to 37.6% compared to the baseline versions of these PNoCs, thereby significantly enhancing reliability, at the cost of up to 19.2% energy overhead and 1.7% photonic area overhead.

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

Keywords: Photonic network on chip; crosstalk noise; reliability

1. INTRODUCTION

Recent developments in silicon photonics have enabled the integration of photonic components with CMOS circuits on a chip. Several photonic network-on-chip (PNoC) architectures have been proposed to date, e.g., [1]-[3], [37], [38]. These PNoCs employ multiple on-chip photonic links that use microring resonator (MR) modulators to convert (i.e., modulate) electrical signals to photonic signals that travel through a photonic waveguide, and MR filter detectors that detect and drop photonic signals on to a photodetector to recover an electrical signal. Each MR has a unique set of (resonance) wavelengths that it can couple to and work correctly with. Typically, photonic waveguides are designed to support dense wavelength division multiplexing (DWDM), where different wavelengths are multiplexed in the waveguide.

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from Permissions@acm.org.

CODES/ISSN '16, October 01-07, 2016, Pittsburgh, PA, USA
© 2016 ACM. ISBN 978-1-4503-4483-8/16/10\$15.00
DOI: <http://dx.doi.org/10.1145/2968456.2968468>

The use of multiple MRs that are in resonance with these wavelengths enables high-bandwidth parallel transfers in photonic waveguides (links).

One of the key challenges for the widespread adoption of DWDM-based PNoC architectures is to minimize the crosstalk noise that emanates from their MRs. In a DWDM-based PNoC, crosstalk occurs when, due to the non-idealities in the resonance of MRs, some optical noise power at an unwanted phase/wavelength mixes with signal power. Crosstalk is an intrinsic property of every MR, so both MR modulators and MR detectors are susceptible to it. The crosstalk noise of MRs negatively affects signal-to-noise ratio (SNR) in waveguides. For instance, the damaging impact of crosstalk noise in the Corona PNoC is presented in [4], where worst-case SNR is estimated to be 14dB, which is insufficient for reliable data transfers, as its corresponding bit-error-rate (BER) is very high, in the order of 10^{-3} .

Crosstalk noise can be classified broadly as homodyne or heterodyne [5]. The *homodyne* crosstalk noise power of a particular wavelength affects the signal power of the same wavelength, whereas in case of *heterodyne* crosstalk, the signal power gets affected by some noise power of one or more other different wavelengths. Both homodyne and heterodyne crosstalk noise may affect signal integrity in a waveguide. Several techniques have been proposed in prior works, e.g., [5]-[7], to mitigate heterodyne crosstalk noise in PNoCs. In general, the sources and effects of heterodyne crosstalk noise in PNoCs are predictable and can be easily controlled by filtering. This fact makes mitigation of heterodyne crosstalk in PNoCs relatively less critical. In contrast, *homodyne crosstalk noise cannot be easily filtered out, as both the noise power and signal power are of the same wavelength*. Moreover, the effect of homodyne crosstalk noise on signal integrity depends on the phase of the noise relative to the signal, which makes its analysis and mitigation more complex. Although, the homodyne crosstalk phenomenon is well understood in relation to WDM long-haul optical networks [8]-[10], it is relatively less explored and understood in relation to PNoCs. Only two prior works analyze the effects of homodyne crosstalk on SNR of PNoCs [4][5]. However, these works neither explain causes of homodyne crosstalk at the device-level, nor present any technique to mitigate it.

Main Contributions: In this paper, we first demonstrate that due to the resonant nature of MRs the circulating power in MRs builds up with time and ultimately recouples back in the waveguide to induce homodyne crosstalk noise. Then, we present a detailed analysis of the characteristics of circulating power in MRs, before presenting a simple, low overhead device-level solution to mitigate the effects of homodyne crosstalk on signal integrity. Our technique is easily implementable in any existing DWDM-based PNoC without requiring major modifications to the architecture. To the best of our knowledge, this is the first work that attempts

to mitigate homodyne crosstalk noise in PNoCs. Our novel contributions in this work are summarized below:

- We demonstrate how the circulating power of an MR, built up due to the resonant nature of the MR, induces homodyne crosstalk noise in the waveguides of a PNoC;
- We perform in-depth analysis and characterization of time-dependent characteristics of circulating power in MRs to understand the behavior of the induced homodyne crosstalk;
- Based on this analysis, we propose a low-overhead homodyne crosstalk mitigation (*HCTM*) technique to control the emanation of homodyne crosstalk noise from MRs by tuning the decoupling of circulating power from the MRs;
- We evaluate the effectiveness and overhead of our *HCTM* technique by implementing it for well-known PNoC architectures, including Corona [1], Firefly [2] and Flexishare [3], running real-world multi-threaded PARSEC benchmarks.

2. BACKGROUND AND RELATED WORK

DWDM-based PNoCs utilize several photonic devices such as microring resonators (MRs) as modulators, switches/injectors, and detectors; photonic waveguides; splitters; and trans-impedance amplifiers (TIAs). The interested reader is directed to [7] for more detailed discussion on each of these components.

As described in [39], an important characteristic of photonic signal transmission in on-chip photonic waveguides is that it is inherently lossy, i.e., the light signal is subject to various types of losses such as through-loss in MR modulators and detectors, modulating losses in modulator MRs, detection loss in detector MRs, propagation and bending loss in waveguides, and splitting loss in splitters. Such losses reduce SNR in photonic waveguides. In addition to the photonic signal loss, crosstalk noise in MRs also deteriorates overall SNR. Both modulators and detectors are susceptible to crosstalk noise in DWDM-based PNoCs. The degraded SNR deteriorates bit-error-rate (BER), jeopardizing the reliability of on-chip data transfers.

As mentioned earlier, crosstalk noise can be classified as homodyne or heterodyne [5]. In case of heterodyne crosstalk, as the noise power and signal power have different wavelengths, the noise power always deteriorates the signal integrity and BER irrespective of its relative phase to the signal. The strength of the heterodyne crosstalk noise at a detector/modulator MR depends on the following four attributes: (i) channel gap between the MR resonant wavelength and the adjacent wavelengths, (ii) Q-factors of neighboring MRs, (iii) data-rate at which the MRs operate [11], and (iv) relative data occurrences in neighboring wavelength-channels [7]. With increase in DWDM, the channel gap between two adjacent wavelengths decreases, which in turn increases heterodyne crosstalk in MRs. With decrease in Q-factors of MRs, the widths of the resonant passbands of MRs increases. Due to the wider passbands of MRs, passband overlap among neighboring MRs increases, which in turn also increases heterodyne crosstalk.

Several techniques have been proposed in prior works, e.g., [5]-[7], [12]-[16], to mitigate heterodyne crosstalk noise in PNoCs. Among them, in [12] and [13], the authors use MR based high-order filters to filter out heterodyne noise from the signal before detection. Xie et al. in [6] analyze the worst-case crosstalk noise and SNR in mesh-based PNoCs, and propose changes in the design of the photonic router to reduce worst-case crosstalk noise. In [7], Chittamuru et al. propose an encoding based approach that reduces undesirable data value occurrences in photonic waveguides to reduce the passband overlap among neighboring MRs, thereby

reducing heterodyne crosstalk noise. Some other prior works also explore heterodyne crosstalk in PNoCs, e.g., [14]-[16], but each of them presents detailed analysis of heterodyne crosstalk noise for a specific PNoC architecture without aiming to mitigate the noise.

For homodyne crosstalk, as the noise power and signal power have the same wavelength, its effects on signal integrity depend on the phase of the noise power. If the homodyne crosstalk noise is in phase with the signal power, the noise power increases the signal power, *increasing* SNR. In contrast, if the homodyne crosstalk noise is out of phase with the signal power, it deteriorates the signal, *decreasing* SNR. Out-of-phase homodyne crosstalk noise, whether coherent (phase-correlated) or incoherent (phase-uncorrelated) to the signal, always degrades signal integrity [4].

The homodyne crosstalk phenomenon in relation to WDM long-haul optical networks is reported in several prior works [8]-[10]. Y. Shen et al. in [10] develop analytical expressions that captures the impact of homodyne crosstalk noise (coherent and incoherent) on an optical signal passing through optical cross-connect nodes in WDM long-haul optical networks. They do not present any noise mitigation technique. Homodyne crosstalk noise has been shown to originate in multiplexer/demultiplexer units and cross-connect nodes of long-haul optical networks, and can be mitigated either by using phase scrambling at the receiver end (as discussed in [9]) or by intelligent management of the optical network (as discussed in [8]). The homodyne crosstalk mitigation techniques discussed in [8] and [9] are specific to the origins of crosstalk noise in and the structure and architecture of the underlying long-haul optical networks. Therefore, these techniques cannot be generalized for use at the chip level to mitigate the homodyne crosstalk noise in PNoC architectures.

As mentioned earlier, homodyne crosstalk is relatively less explored and understood at the chip level and in relation to PNoC architectures. Only two prior works (i.e., [4] and [5]) have shown homodyne crosstalk noise to originate from MRs of PNoCs and have analyzed the effects of homodyne crosstalk on the SNR of PNoCs. Nikdast et al. in [5] analyze the worst-case homodyne incoherent crosstalk noise and resulting SNR in arbitrary fat-tree-based PNoCs. Duong et al. [4] analyze homodyne coherent crosstalk noise in ring-based DWDM PNoCs. As demonstrated in [4], the worst-case SNR due to homodyne coherent noise in the data waveguide of Corona PNoC is estimated to be 14dB, which corresponds to BER of 10^{-3} . Now, as explained in [6], the value of BER should be 10^{-9} or less for reliable communication, which implies that the BER of 10^{-3} obtained in [4] for the Corona PNoC can severely harm the communication reliability. This implies that, in PNoCs, homodyne noise can be severe and should be significantly mitigated to ensure reliable communication. *But none of the prior works [4] and [5] aims to mitigate homodyne crosstalk noise.*

Now, one way of mitigating the reliability concerns related to the homodyne crosstalk noise is using error-correcting codes (ECCs). However, ECCs cannot eliminate all the faults or errors in a PNoC, when the BER is very high in the order of 10^{-3} or higher. Moreover, ECCs, whether implemented in software or hardware, consume high energy and incur significantly high performance overhead. In addition to this high overhead, ECCs do not guarantee elimination of all errors. Thus, the effectiveness of ECCs in mitigating the homodyne crosstalk noise is questionable. On the other hand, any higher-level homodyne noise mitigation technique, other than ECCs, depends on the system's ability to differentiate between the noise power and the signal power. But as the wavelength of homodyne noise power is same as that of the signal power, it is very difficult to differentiate or filter the noise

power from the signal power at any higher level. Therefore, the effectiveness of such higher-level techniques is also questionable.

In contrast to the ECCs and other high-level techniques, in this paper, we present a low-overhead device-level technique called HCTM to mitigate homodyne crosstalk. Our HCTM technique extenuates the root cause of homodyne crosstalk noise, i.e., the circulating power in MRs. This enables HCTM to eliminate all the errors incurred due to the homodyne crosstalk noise. Thus, it can be implied that our HCTM technique is better than ECCs and any such higher-level technique in ensuring the reliability of communication in PNoCs.

3. HOMODYNE CROSSTALK: CAUSE/EFFECTS

In Section 3.1, we present the general properties of MRs. Then we present detailed analysis of spectral and temporal characteristics of MRs in Section 3.2 to understand homodyne crosstalk in PNoCs. Finally, in Section 3.3, we discuss modeling of homodyne crosstalk in PNoCs.

3.1 General Properties of MRs

Before going into the specifics of homodyne crosstalk noise, we first discuss the general properties of MRs. Optical MRs are extensively described in literature [17][18]. As shown in Fig. 1(a), an MR consists of a looped optical waveguide coupled to a straight bus waveguide (BWG). The straight BWG acts as a coupling mechanism to access the loop. The looped waveguide of the MR could be circular in shape (Fig. 1(a)) or elongated (Fig. 1(b)). The elongated MR with a straight section along one direction (typically along the coupling section) is called *racetrack* MR. The elongated shape of the racetrack MR renders better control of the coupled optical power [18]. Otherwise, the racetrack MR is functionally and characteristically similar to the circular MR. Therefore, the analysis presented in this section is equally applicable to both types of MRs. Moreover, the analytical models presented in this section and subsequent sections hold true for both types of MRs and can be used to capture the crosstalk noise characteristics of both types of MRs.




Fig. 1: Illustrations of waveguide coupled photonic microring resonators (MRs): (a) circular MR (b) racetrack MR.

As shown in Fig. 1(a), the BWG that is coupled to the MR has two ports termed as *input* and *pass*. The coupled optical field (E_c), the circulating optical field (E_{circ}), and the output optical field (E_{pass}) in the figure depend on the input optical field (E_{in}) through the field transmission coefficient (t) and field-coupling coefficient (κ). The relationship among these optical fields can be given by the following equations, which are derived from the transfer matrix model described in [18]:

$$E_c = j\kappa E_{in} + t E_{circ}, \quad (1)$$

$$E_{pass} = E_{tr} + E_N = t E_{in} + j\kappa E_{circ}, \quad (2)$$

$$E_{circ} = a e^{j\varphi} E_c, \quad (3)$$

Here, $j = \sqrt{-1}$, which represents π phase shift of the ideal MR to BWG coupler, a is the roundtrip amplitude transmission coefficient, including both propagation loss in the ring and loss in the couplers, and φ is roundtrip phase shift of the field. The coefficients t and κ are such that $t^2 + \kappa^2$ are the power splitting ratios of the coupler, and they are assumed to satisfy $t^2 + \kappa^2 = 1$, which means there are no extra losses in the coupling section. As evident from Eq. (1)-(3), E_{circ} provides a positive feedback for E_c , which ultimately affects E_{circ} and E_{pass} . As a result, in the presence of non-zero E_{in} , E_{circ} grows with every roundtrip due to increasing E_c , resulting in an increase of circulating optical power (P_{circ}) with time. From Eq. (2), a part of E_{circ} recouples back in the BWG from the MR as a noise field ($E_N = j\kappa E_{circ}$), which adds up with the transmitted field ($E_{tr} = t E_{in}$) to give E_{pass} as the output field. In fact, E_N is the homodyne noise field, as its wavelength is the same as the wavelength of E_{tr} and E_{in} .

For a given cross-section of the BWG and the MR's loop waveguide, the optical power travelling in these waveguides is proportional to the square of the traversing optical field [17]. This implies that the relationship between the coupled optical power (P_c), circulating optical power (P_{circ}), and output optical power (P_{pass}) to the input optical power (P_{in}) can be derived from Eq. (1)-(3), which is given in the following equations [17]:

$$B = \frac{P_{circ}}{P_{in}} = \left| \frac{E_{circ}}{E_{in}} \right|^2 = \frac{\kappa^2 a^2}{1 - 2at \cos \varphi + a^2 t^2}, \quad (4)$$

$$P_{pass} = P_{tr} + P_N = \frac{a^2 - 2at \cos \varphi + t^2}{1 - 2at \cos \varphi + a^2 t^2} P_{in}, \quad (5)$$

$$P_N = - \left(\frac{2at\kappa^2 e^{j\varphi}}{1 - at e^{j\varphi}} + B\kappa^2 \right) P_{in}, \quad (6)$$

$$P_{tr} = t^2 P_{in}, \quad (7)$$

Here, B is the buildup factor, which is the ratio of P_{circ} to P_{in} (Eq. (4)). Eq. (5) gives the value of P_{pass} in terms of t , a , and φ . This equation can be broken down in Eq. (6) and Eq. (7) to show two components of P_{pass} : P_N (noise power) and P_{tr} (transmitted signal power). Note that P_N (Eq. (6)) is homodyne noise power, as its wavelength is the same as that of P_{tr} and P_{in} .

Note that the noise characteristics of MRs that are captured by the models presented in Eq. (1)-(7) do not differ for circular and racetrack MRs. Only the amount of noise power in MRs differs between the two types of MRs. It is because only the values of coefficients a , t and κ differ between circular and racetrack MRs.

In summary, for a non-zero input power (P_{in}), the circulating power (P_{circ}) in the MR builds up by the factor B at steady state, which ultimately contributes to the homodyne noise power (P_N) in the BWG (factor $B\kappa^2$ in Eq. (6)).




Fig. 2: Illustration of gradual increase and decrease of power circulating in MR: (a) when $P_{in} \neq 0$ (b) when $P_{in} = 0$.

3.2 Spectral and Temporal Characteristics of P_N

As evident from Eq. (6), the homodyne noise power (P_N) depends on roundtrip phase shift φ . The value of φ changes periodically with wavelength of operation, which implies that use of Eq.

(5), (6) requires determination of φ that is equal to $2\pi Ln_{eff}/\lambda$, where L is the roundtrip length of MR and λ is operating wavelength. This requires complicated calculation of n_{eff} that depends on MR geometry and materials. A more commonly used equation for the power transmission (and hence for the noise power) on and off resonance of an MR is the Lorentzian transfer function [19]:

$$T_N = \frac{(P_N^{\varphi=2\pi}/P_{in})}{1 + (2Q(\Delta\lambda/\lambda_r))^2}, \quad (8)$$

where, Q is loaded quality factor of the MR, λ_r is MR's resonance wavelength, and $\Delta\lambda$ is the difference between the operating wavelength and λ_r . Note that, for $\Delta\lambda=0$, Eq. (8) gives normalized homodyne crosstalk noise (normalized to P_{in}) for on-resonance MR operations, whereas for $\Delta\lambda \neq 0$, Eq. (8) gives normalized homodyne crosstalk noise for off-resonance MR operations.

Note that, at the *critical coupling condition* (when $t=a$ in Eq. (5)), whether on-resonance or off-resonance, P_{pass} is minimized, but the effect of P_N on P_{tr} and hence on P_{pass} is maximized. This implies that the critical coupling condition does not alleviate P_N , even though it minimizes P_{pass} .

In our analysis so far, we have considered only the spectral properties of P_N , which characterizes the behavior of P_N subject to various wavelengths of light. As evident from Eq. (6), in addition to the parameters φ , t , a , and κ , the noise power P_N also depends on the P_{circ} build-up factor B and input power P_{in} . Parameters B and P_{in} not only have spectral properties (such as operational wavelength, field phase φ), but also have temporal properties, i.e., B and P_{in} may vary with time. For instance, in the PNoC paradigm, the value of P_{in} typically keeps flipping between logic '1' and logic '0' with time.

As mentioned earlier, in the presence of a non-zero P_{in} , P_{circ} in an MR gradually increases with every roundtrip (and hence with time) and builds up by the factor B at steady state. As P_N depends on P_{in} and P_{circ} ($B\kappa^2 P_{in}$ - Eq. (6)), P_N also follows the same trend as P_{circ} . This fact is depicted in Fig. 2(a), which illustrates how P_{circ} and P_N grow with every roundtrip in the presence of a non-zero P_{in} . If P_{in} drops to zero while P_{circ} is growing, then P_{circ} starts decaying with time from whatever value it had grown to before P_{in} dropped to zero. This fact is depicted in Fig. 2(b), which illustrates how P_{circ} and P_N decays with every roundtrip when $P_{in}=0$. The steady state value of P_{circ} (and hence of P_N) and the rate with which P_N grows/decays every round trip (or with time) depends on the cavity photon lifetime (τ_p) and Q -factor of the MR through a natural exponential function. The growth/decay of P_N as a function of time is expressed as [20]:

$$P_N(t) = \begin{cases} T_N (1 - e^{-(t/\tau_p)}), & P_{in} \neq 0 \\ T_N^{t=0} (e^{-(t/\tau_p)}), & P_{in} = 0 \end{cases}, \quad (9)$$

$$\tau_p = Q\lambda/2\pi c, \quad (10)$$

where, T_N can be obtained from Eq. (8). When $P_{in} \neq 0$, P_N grows exponentially until it reaches the steady state value T_N in τ_p time. When $P_{in}=0$, P_N decays exponentially starting from its initial value of $T_N^{t=0}$. As the value of P_{in} typically keeps flipping between logic '1' and logic '0' with time in a typical PNoC waveguide, it implies that P_{in} consists of a sequence of falling and rising edges. Therefore, it can be inferred that the value of P_N at any given time depends on the instantaneous value of P_{in} and on the amount of time that has elapsed from the last falling or rising edge of P_{in} . Moreover, the time between two successive falling/rising edges of P_{in}

depend on the bit-period and the bit-pattern of P_{in} , which implies that P_N also depends on bit-period (bit-rate) and bit-pattern of P_{in} .

In summary, the homodyne crosstalk noise (P_N) originating from MRs varies with field phase φ , which in turn depends on the operating wavelength. P_N also varies with time, and the time-dependent variation profile of P_N depends on the bit-period (and hence the bit-rate) and bit-pattern of the input signal (P_{in}) in the photonic waveguide, along with some time-independent properties of MRs such as photon lifetime (τ_p) and Q . The time-varying homodyne noise (P_N), when coupled in the BWG, gets mixed up with the time-varying transmitted signal power (P_{tr}) causing random fluctuations in the amplitude of the output signal power (P_{pass}). These random fluctuations in P_{pass} reduce the optical eye opening, which in turn increases the BER and reduces the communication reliability in PNoCs.

3.3 Modeling of Homodyne Noise (P_N) in PNoCs

Our ultimate goal in this paper is to analyze and mitigate the homodyne crosstalk noise in PNoC architectures. As mentioned earlier, PNoCs are typically designed to support DWDM, where a large number of wavelengths are multiplexed in a single BWG. The group of multiple MRs that are in resonance with these DWDM wavelengths and that are "arrayed" along the BWG can be referred to as an *MR bank*. This MR bank can be a modulator bank or detector bank. In a typical PNoC, the BWG passes through multiple nodes, and a pair of such modulator and detector MR banks is used at every node, which enables high bandwidth parallel data transfers in the BWG. Thus, modulator banks, detector banks, and BWGs are basic building blocks of a DWDM PNoC. Before we extend our homodyne noise model (discussed in previous subsections) to the entire PNoC, it is important to first extend it to these basic building blocks. Our aim is to reduce the worst case P_N in a PNoC, thus we model only the worst-case value of P_N in MR banks, BWG, and the PNoC. Note that the worst-case P_N occurs at steady state ($t=\tau_p$ in Eq. (9)), thus, the value of worst-case P_N does not change with time.

Modeling of P_N in MR Modulator and MR Detector Banks:

An MR modulator bank at a node along the BWG of a PNoC could be in either active state or inactive state. Similarly, an MR detector bank could be in either active state or inactive state. All MRs in an active modulator bank become ON and OFF resonance with their corresponding wavelengths with time to modulate a particular bit pattern on their corresponding wavelengths. All MRs in an inactive modulator bank and in an inactive detector bank are OFF resonance with their corresponding wavelengths, whereas the MRs of an active detector bank are ON resonance with their corresponding wavelengths. The worst-case homodyne crosstalk noise for a DWDM wavelength emanating from a modulator or a detector MR bank can be modeled by extending Eq. (8) as follows:

$$P_N^{Bank}(\lambda_j) = \sum_{i=0}^{n-1} \frac{P_N^{\varphi=2\pi}/P_{in}}{1 + (2Q\{|\lambda_i - \lambda_j|/\lambda_i\})^2}, \quad (11)$$

where, n is the number of MRs in the bank that is the same as the number of DWDM wavelengths. $P_N^{\varphi=2\pi}$ is calculated by putting $\varphi=2\pi$ in Eq. (6). λ_j is the wavelength of P_N . λ_i is the resonance wavelength of the i^{th} MR of the bank. In case of an active modulator bank or detector bank, for a given value of λ_j , there exists exactly one i (and corresponding MR) for which $\lambda_j = \lambda_i$. In contrast, in case of inactive modulator bank or detecting bank, for a given value of λ_j , there exists no i for which $\lambda_j = \lambda_i$. Note that λ_j gets picked up by each MR in the bank, the extent of which depends on

the distance $|\lambda_i - \lambda_j|$ of λ_j from the resonance wavelength (λ_i) of the corresponding MR. Therefore, optical power of wavelength λ_j builds up in each MR of the bank, which ultimately contributes to the homodyne crosstalk noise of wavelength λ_j . Thus, Eq. (11) models the total worst-case homodyne crosstalk noise emanating from all the MRs of the bank combined.

Modeling of P_N in DWDM Bus Waveguide (BWG):

When a source node sends a data packet along a BWG of a PNoC to a destination node, the modulator bank of the source node remains active, the detector bank of the destination node remains active, and all other modulator and detector banks along the BWG remain inactive. In this scenario, all wavelengths of the DWDM spectrum are modulated by the source modulator bank and then they travel along the BWG from the source node to the destination node where they are detected by the destination detector bank. For a particular wavelength from the DWDM spectrum, the homodyne crosstalk noise at that wavelength originates from all the intermediate MR banks along the BWG that add up and contribute to the total noise power $P_N^{BWG}(\lambda_j)$ at the destination detector MR. The worst-case homodyne crosstalk noise at a wavelength in a DWDM bus waveguide between the source modulator bank and the destination detector bank can be assessed using the following equation:

$$P_N^{BWG}(\lambda_j) = \sum_{l=1}^{L_D+L_M} P_N^{Bank}(j) + \sum_{m=1}^{M_D+M_M} P_N^{Bank}(j), \quad (12)$$

Here, L_D and L_M are the numbers of active MR detector banks and active MR modulator banks respectively, whereas M_D and M_M are the numbers of inactive MR detector banks and inactive MR modulator banks, respectively.

4. MITIGATION OF HOMODYNE CROSSTALK

As implied from the conclusions of our analysis presented in the previous section and Eq. (6), the worst-case homodyne crosstalk noise P_N originating from a particular MR depends on the MR's power build-up factor B and power-coupling coefficient κ^2 . As evident from Eq. (4), B is a function of the roundtrip power-transmission coefficient a^2 . Therefore, P_N can be mitigated by reducing either κ^2 or B (hence a^2). However, reducing κ^2 increases t^2 (as $\kappa^2 = 1 - t^2$), which in turn increases P_{tr} (Eq. (7)). The increase in P_{tr} reduces the extinction ratio [17], which increases the susceptibility of signal transmission (P_{pass}) to noise (P_N) as evident from Eq. (5). Therefore, reducing κ^2 is not a desirable option for mitigating P_N . Alternatively, reducing a^2 can be achieved by increasing the roundtrip power loss ($1 - a^2$) in an MR. The roundtrip loss in an MR is in general attributed to propagation loss and bending loss in an MR's looped waveguide, and coupling loss at the MR-BWG coupling region [17]. The propagation loss can be increased by increasing either the sidewall roughness of or the free carrier concentration in the MR's waveguide [21]. However, increasing free carriers would alter the resonance of the MR, and increasing sidewall roughness is difficult to control.

We propose a simple and elegant method of increasing loss in an MR to reduce homodyne P_N with more predictable control. Our homodyne crosstalk mitigation (HCTM) technique has two components. First, we propose to use racetrack MRs (Fig. 3), which adds excess bend losses in MRs due to mismatch losses at the straight-bend transition of MRs [17]. Second, as shown in Fig. 3, we introduce a secondary coupling waveguide for an MR that can decouple most of the circulating power (P_{circ}) out of the MR before it can couple back in the BWG, thus increasing loss ($1 - a^2$) to reduce B , and noise power P_N (as per Eq. (4), (6)).

Decoupling P_{circ} out of the MR greatly reduces P_N irrespective of the temporal variation profile of P_{in} . Thus, our HCTM technique mitigates the homodyne noise in PNoC architectures by extenuating its root cause, i.e., P_{circ} in MRs. Therefore, it can be implied that HCTM technique is agnostic to the underlying PNoC architecture, as it reduces homodyne crosstalk noise in all types of PNoC architectures in the same way, i.e., by extenuating P_{circ} in MRs.




Fig. 3: Illustration of decoupling waveguide (DWG) implementation for: (a) modulator MRs, (b) detector MRs, (c) switch MRs.

We refer to this secondary coupling waveguide as decoupling waveguide (DWG). The inclusion of the DWG is functionally similar to using add-drop port MRs (described in [18]) except that the DWG has an absorbent/grating structure at its drop port, which absorbs or scatters out all the power that is dropped at the drop port. The absorbent/grating structure is the same as used in [22] as part of the dithering signal based mechanism proposed for in-situ, high-speed detection of temperature and process variation in silicon photonic chips. Note that an increase in round-trip loss decreases MR's Q factor, which may result in increase of heterodyne crosstalk, as discussed in Section 2. The decrease in Q may also increase the through loss of the MR, resulting in a decrease of signal strength and consequently a decrease in SNR. This increase in heterodyne crosstalk noise and MR's through loss can be mitigated by using higher-order MR filters with detector MRs [12][13], in conjunction with encoding based techniques [7]. The use of higher-order filters with detector MRs is illustrated in Fig. 3(b), the details of which are presented later in this section. Moreover, note that the introduction of a DWG along the BWG incurs negligible amount of inter-waveguide crosstalk, as we assume minimum $10\mu\text{m}$ separation (equal to the diameter of the coupled racetrack MR) between the DWG and BWG.

As shown in Fig. 3(c), a DWG can be used to decouple P_{circ} from a switch MR as well. Optical switch MRs are typically coupled into two waveguides. As a result, they have four ports (Fig. 3(c)). As a switch MR is already coupled to two parallel waveguides, the introduction of a third waveguide as DWG in the same plane would incur extra loss due to waveguide crossings. Therefore, to avoid waveguide crossings and related losses, we introduce the DWG for a switch MR on a separate layer below or above the MR. A monolithic integration of DWG on a separate layer above or below the layer of other photonic devices is possible with the monolithic multilayer integration technology described in [33]. As discussed in [33] and [36], this type of monolithic multilayer integration of photonic devices using CMOS back-end compatible materials can be highly area-efficient, energy-efficient and cost-effective. In the absence of a DWG, the circulating power in an inactive/active (OFF-state/ON-state) switch MR can

recouple back at the pass/output port, which will induce homodyne crosstalk at the pass/output port. Nevertheless, in the presence of a DWG, P_{circ} is decoupled from the MR to the DWG, which mitigates the homodyne crosstalk noise at pass/output port of the switch MR. Fig. 3(b) shows two cascaded racetrack rings, which manifests a symbol representing a higher order filter for the embedded Si-Ge detector. If this detector MR is in an inactive state, P_{circ} needs to be decoupled from it to reduce the homodyne crosstalk noise. In contrast, if P_{circ} is decoupled when the detector MR is in active state, then the signal strength (which in fact depends on P_{circ} in this case) reduces, which in turn deteriorates the SNR.

The above discussion implies that a mechanism is needed for the DWG to detune its decoupling capability when the detector MR is in active state. This motivates the design of a tunable DWG (TDWG), which is described in detail in the next section.

5. TUNABLE DECOUPLING WAVEGUIDE

Before we describe the design of a TDWG, it is important to understand how a TDWG works. The straight section of a race-track MR, which is parallel to the TDWG (Fig. 4), works as a directional coupler. Some part of the optical power that travels along this straight section of the MR is coupled to the TDWG depending on the design of the TDWG. Typically, a TDWG is a standard Si-SiO₂ waveguide, with a Si core and SiO₂ cladding. The coupling of power between the MR and the TDWG is given by the following equations [24]:

$$P_{TDWG}(z) = P_{MR} \frac{K_d^2}{K_d^2 + \delta^2} \left[\sin \left\{ (K_d^2 + \delta^2)^{1/2} z \right\} \right]^2, \quad (13)$$

$$\delta = (\beta_{MR} - \beta_{TDWG})/2, \quad (14)$$

$$P_{MR} = (a^2 \kappa^2 P_{in})/2, \quad (15)$$

Here, K_d is coupling coefficient, P_{TDWG} is the amount of power decoupled from the MR into the TDWG, β_{MR} and β_{TDWG} are propagation constants of light in the MR waveguide and TDWG respectively, z is the length of the straight section of the MR, and P_{MR} is the optical power in the MR waveguide at the coupling region. From Eq. (13)-(15), P_{TDWG} depends on K_d . As discussed in [24], K_d depends on field penetration depth (d), gap (g), and the refractive index of the MR waveguide's core and TDWG's core (n_{Si}). This implies that K_d , and hence P_{TDWG} can be controlled by tuning n_{Si} . The change in n_{Si} can be achieved by changing the TDWG's Si core free carrier concentration, which works on the same principles of free carrier dispersion as in the state-of-the-art MR modulators [21]. Similar to MR modulators, the TDWG can be doped like a PN-junction to function in a reverse-biased manner to render better control on free carrier concentration in the TDWG's Si core with faster response.




Fig. 4: Cross-sectional structure of tunable decoupling waveguide (TDWG): (a) OFF-state TDWG (b) ON-state TDWG.

Fig. 4 depicts the cross-sectional structure of the TDWG's PN-junction along with the gap g and penetration depth (coupling depth) d . Fig. 4(a) and (b) illustrate how the decoupling capabilities of a TDWG can be tuned ON and OFF by changing the reverse bias condition of the constituent PN-junction.

As shown in Fig. 4(a), at zero bias condition, the PN-junction has a narrower depletion region resulting in higher free carrier concentration in the core. In this case, the coupling depth (d) of the TDWG is smaller than the gap g . In contrast, as shown in Fig. 4(b), an application of a non-zero reverse bias across the junction increases the depletion width resulting in the decrease of carrier concentration, which in turn increases n_{Si} causing d to penetrate beyond g . When d penetrates beyond g , a significant amount of power is decoupled from the MR into the TDWG core. Thus, n_{Si} of a TDWG can be altered by changing the free carrier concentration in the TDWG core, which ultimately tunes the TDWG decoupling capability between ON and OFF states.

It is imperative to design n_{Si} and carrier concentration values of the PN-junction appropriately to get the right amount of decoupling for ON and OFF states of the TDWG. To understand the relationship among the decoupled power P_{TDWG} , free carrier concentration ($N_e=N_h$), and n_{Si} , we plot P_{TDWG} (in dB) and $N_e=N_h$ versus n_{Si} in Fig. 5 for a TDWG. We sweep n_{Si} from 2.4 to 4.1. As shown in Fig. 5, the TDWG can be designed to decouple up to 30dB power from the MR. The smallest value of decoupled power (0.1dB) can be achieved for $n_{Si}=2.4$ and $N_e=6 \times 10^{20}$. At this low value of decoupled power, the TDWG can be considered in OFF state, as it decouples negligible power from the MR.




Fig. 5: Decoupled power and free-carrier concentration (N_e) vs. n_{Si} of TDWG for ACTIVE, INACTIVE state MRs. ($g=200\text{nm}$, $z=2\mu\text{m}$).

We propose to dope the TDWG's PN-junction (Fig. 4(a)) to achieve $N_e=6 \times 10^{20}$ at zero bias condition. Thus, the default state of the TDWG is OFF state. As evident from Fig. 5, the TDWG can be switched ON by applying a reverse biased voltage across the junction, which would in turn decrease N_e below 6×10^{20} causing decoupled power of greater than 0.1dB. The target value of decoupled power for the ON-state TDWG should be chosen based on the underlying BWG and PNoC architectures, so that the worst-case P_N emanated from MRs is minimized by maximizing the extenuation of P_{circ} from MRs. For a given value of signal power, the worst-case (maximum) P_N results in the worst-case (minimum) SNR. Moreover, the value of decoupled power that maximizes the extenuation of P_{circ} not only minimizes the worst-case P_N , but also guarantees the minimization of the best-case and average-case P_N . Therefore, the minimization of the worst-case P_N not only maximizes the worst-case SNR, but also guarantees the maximization of the best-case and average-case SNR values.

However, note that the maximization of the worst-case SNR is more important than the maximization of the best-case or average-case SNR, as the worst-case SNR determines the BER, which in

turn affects the reliability of communication in PNoCs. Therefore, the target value of decoupled power for the ON-state TDWG should be chosen based on the underlying BWG and PNoC architectures, so that the worst-case SNR is maximized. As an illustration of this fact, we present an architecture-specific implementation of TDWG for the Corona PNoC in Section 5.1. Note that the use of TDWG to mitigate homodyne crosstalk incurs overhead, the analysis of which is presented in Section 5.2.




Fig. 6: Demonstration of (a) worst-case homodyne noise in Corona PNoC's data bus waveguide; (b) use of tunable decoupling waveguide in Corona PNoC's data bus to mitigate homodyne noise.

5.1 Implementation of TDWG for the Corona PNoC

We now demonstrate how the use of TDWG can mitigate homodyne crosstalk noise in the Corona PNoC. We direct the reader to [1] for detailed information about the Corona PNoC architecture. Briefly, Corona consists of 256 general-purpose cores grouped into 64 four-core clusters. These clusters are connected together through three bus waveguides (BWGs) that include an optical crossbar for data communication, a broadcast bus for multi-casting, and an arbitration waveguide. The main laser source is fed into a loop and split into these BWGs. The optical crossbar for data communication is comprised of a multiple-write-single-read (MWSR) BWG, which starts at cluster 0, passes through 62 intermediate clusters, and ends at cluster 63. Each cluster has a bank of 64 MRs arrayed along the BWG. When an optical signal travels from cluster 0 to cluster 63 along the BWG, all the MRs (total 4096 MRs; 64 clusters having 64 MRs each) along the BWG incur MR through losses in the signal. Thus, a very large number of utilized MRs incur a very high value of optical signal loss in Corona data BWG, which results in very poor signal strength. A very poor signal strength in the Corona data BWG makes it more susceptible to homodyne noise.

Fig. 6(a) shows the worst-case homodyne noise scenario in the data BWG. From the figure, the modulator bank of cluster 0 is in active state and it modulates a data packet on the DWDM wavelengths of the BWG. This data packet travels along the BWG and passes through inactive-state MR modulator banks of 62 intermediate clusters before reaching the active-state MR detector bank of cluster 63. As shown in the figure, while the data packet travels through cluster 0 to cluster 63, the modulator banks of all 62 intermediate clusters incur homodyne crosstalk noise in the data packet. This noise is picked up by the detector bank of cluster 63 along with the modulated data packet.

Fig. 6(b) demonstrates the use of TDWGs to mitigate the homodyne crosstalk noise in the Corona data BWG. As shown in

the figure, the TDWGs corresponding to the modulator banks of cluster 0 to cluster 63 are in ON state, whereas the TDWGs corresponding to the detector bank of cluster 63 are in OFF state. The ON-state TDWGs of intermediate modulator banks decouple the circulating power from their respective MRs, which in turn greatly reduces the homodyne crosstalk noise. As discussed in Section 4, we propose to use higher-order (3rd order or higher) filters with detector MRs in the detector bank to reduce the heterodyne crosstalk noise, which is illustrated in the figure.

To select the most appropriate value of decoupled power for ON-state TDWGs, we evaluated the worst-case SNR in the data BWG and the control BWG (described in [1][4]) of the Corona PNoC as a function of decoupled power using the equations in Section 3 and 5. In our calculation of SNR, we also factored in the insertion losses incurred by the splitters, waveguide, and MRs, along with the heterodyne crosstalk noise incurred due to the passband overlap of MRs [34][35]. The plots of the worst-case SNR are shown in Fig. 7. From the figure, the improvements in SNR for the data BWG and the control BWG saturate for values of decoupled power greater than 15dB and 20dB respectively. Therefore, we select 15dB as the target value of decoupled power for TDWGs in the Corona data BWG and 20dB for the control BWG. The architectures of the Corona data BWG and control BWG differ from each other, which results in different values of desired decoupled power for them. Hence, it can be concluded that the target values of decoupled power for ON-state TDWGs should be chosen based on the underlying BWG and PNoC architecture.

From Fig. 7, the worst-case SNR for Corona PNoC's data BWG increases from about 11dB to 15dB when the decoupled power is increased from about 0.1dB to 15dB. Now, as discussed in [6], the worst-case SNR, when expressed in absolute value (not in dB), is related to the BER through the following equation: $BER = 0.5 \times \exp(-SNR/4)$. From this relation between the worst-case SNR and BER, the increase in the worst-case SNR from 11dB to 15dB corresponds to the improvement (decrease) in BER from 2.1×10^{-2} to 1.8×10^{-4} . From [6], the maximum allowable BER for reliable communication is 10^{-9} . This implies that even after achieving about $100 \times$ improvement in BER (from 2.1×10^{-2} to 1.8×10^{-4}), obtained by minimizing the worst-case homodyne noise, sufficiently reliable communication cannot be achieved. This is because, even after minimizing the homodyne crosstalk noise, the heterodyne crosstalk noise still exists in PNoCs, which needs to be significantly mitigated to achieve BER in the order of 10^{-9} or less and reliable communication in PNoCs. However, note that this paper focuses only on mitigating the homodyne crosstalk noise, and the problem of heterodyne crosstalk mitigation is out of the scope of this work.




Fig. 7: Worst-case SNR vs. decoupled power for Corona PNoC's data and control waveguides.

5.2 Device-Level Overhead Analysis

We now evaluate the area and energy overhead incurred by our proposed TDWGs. The default state of a TDWG is the OFF state,

where it does not consume static or dynamic energy. A TDWG consumes dynamic switching energy every time it is switched ON, and after switching ON it consumes static power related to the reverse saturation current in the constituent PN-junction for the entire time it is ON. As illustrated in Fig. 3, only the TDWGs corresponding to detector MRs need to dynamically switch ON and OFF. The TDWGs corresponding to modulator and switch MRs always remain ON. Therefore, the TDWGs corresponding to modulator and switch MRs do not consume dynamic switching energy, but consume only static power. In contrast, as the TDWGs corresponding to detector MRs need to switch OFF while the detector MRs are detecting data bits and then switch ON immediately after the detection event, they consume dynamic energy every time a data packet is transferred.

The amount of static power and dynamic energy consumed depends on the level of decoupled power in the ON-state TDWGs. In Fig. 8, we plot static power density (fW/cm²) and dynamic energy consumption values versus decoupled power. From the figure, static power density decreases and dynamic energy increases with increase in decoupled power. Moreover, it can be observed that for the 15dB decoupled power chosen for the Corona data BWG, the values of dynamic energy and static power density are 4.3pJ and 31.76fW/cm² respectively.




Fig. 8: Static power density and dynamic energy overhead of tunable decoupling waveguide versus decoupled power.

Furthermore, as mentioned in Fig. 5, we select the length of the MR's straight section (z) that is coupled to the TDWG to be $2\mu\text{m}$. As depicted in Fig. 4, the width (w) of the TDWG is 450nm . Therefore, if we consider the length of a TDWG to be $2\mu\text{m}$ more to account for its portion that is not coupled to the MR, and if we assume the length and the width of the absorbent/grating to be $1\mu\text{m}$ each, then a single TDWG along with the grating would consume $2.8\mu\text{m}^2$ surface area. As depicted in Fig. 4, the height of a TDWG is 250nm , which results in the volume of a single TDWG along with the grating to be $0.7\mu\text{m}^3$.

Note that the introduction of TDWGs incur extra fabrication cost that is related to the amount of utilized silicon real estate. The amount of utilized silicon real estate is proportional to the surface area and volume overhead of TDWGs. As mentioned earlier, a single TDWG along with the grating would consume $2.8\mu\text{m}^2$ surface area, which is very less compared to the surface area of $78.5\mu\text{m}^2$ consumed by a typical circular MR. This implies that the introduction of TDWGs incur minimal cost overhead related to device fabrication. Moreover, TDWGs have significantly low engineering design cost as well, as the structure of the TDWGs is simple and fairly repetitive, which reduces the complexity of mask-set designs and fabrication process.

6. EVALUATION

6.1 Evaluation Setup

To evaluate our proposed homodyne crosstalk noise mitigation (*HCTM*) approach in DWDM-based PNoCs, we implement and

integrate it with three well-known crossbar-based PNoCs: Corona [1], Flexishare [2], and Firefly [3]. We modeled and performed simulation based analysis of these enhanced PNoCs using a cycle-accurate NoC simulator, for a 256-core system at 22nm. GEM5 full-system simulation [26] of parallelized PARSEC applications [25] 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 $20\text{mm} \times 20\text{mm}$ chip, to determine lengths of waveguides in the Corona, Firefly and Flexishare PNoCs. Based on this analysis, we estimated the time needed for light to travel from the first to the last node as 8 cycles in all of these PNoCs at 5 GHz clock frequency. We use a 512-bit packet size, as advocated in all of these PNoCs.

The static and dynamic energy consumption of electrical routers and concentrators in Corona, Firefly, and Flexishare is based on results from the DSENT [27] tool. We evaluate area and energy overheads for our *HCTM* technique when implemented with the above mentioned PNoCs, the results of which are given in Section 6.3. For energy consumption of photonic devices, we adapt parameters from recent work [28]-[30], with 0.42pJ/bit for every modulation and detection event and 0.18pJ/bit for the driver circuits of modulators and photodetectors. We used 0.02dB , 1dB/cm , $0.005\text{dB}/90^\circ$, and 0.5dB values of MR through loss, waveguide propagation loss, waveguide bending loss and coupling loss respectively, to determine the photonic laser power budget and correspondingly the electrical laser power. The MR trimming power is set to $130\mu\text{W/nm}$ [31] considering a 40°K trimming range.

6.2 Evaluation Results for State-of-the-art PNoCs

First, we evaluate the worst-case SNR in Corona, Flexishare and Firefly PNoCs when used with our *HCTM* technique, and compare the results with the SNR in their respective baselines. We considered two different baseline configurations for each PNoC architecture: 1) a configuration with only heterodyne crosstalk without any homodyne crosstalk (ideal case); 2) a configuration with homodyne as well as heterodyne crosstalk.

In the Corona PNoC with token ring arbitration [1], we consider multiple write single read (MWSR) data BWGs as well as control/arbitration BWGs for our evaluation of worst-case SNR. Likewise, we consider the multiple write multiple read (MWMR) data BWGs of Flexishare PNoC [3] for evaluation. The Flexishare PNoC is a 64-radix, 64-node architecture with 4 cores in each node having 32 data channels for inter-node communication. We also consider the data BWGs of the Firefly PNoC [2], which are configured as reservation-assisted single write multiple reader (R-SWMR) data BWGs.

Generally, in a PNoC, for given crosstalk noise, the worst-case SNR happens at the MR detector bank of a node for which the loss in signal is the highest, as the highest loss of signal renders the worst signal power and hence the worst SNR. We refer to such a node as the worst-case power loss node (WCPLN). Cluster #63 in Corona PNoC and node 63 (R63) in Flexishare PNoC are the WCPLNs. Similarly, in Firefly PNoC [2], the router 0 of cluster 4 (C4Ro) is the WCPLN. We utilize the models presented in [4] (for heterodyne crosstalk) and Section 3 and 5 of this paper (for homodyne crosstalk) to calculate the total received crosstalk noise (including both homodyne and heterodyne crosstalk noise) and SNR at the detectors of these WCPLNs. Note that the worst-case heterodyne crosstalk noise for a data BWG of any PNoC occurs when all the 64-bits of the received data packet are 1's. In contrast, as implied from the discussion given in Section 5.1, the worst-case

homodyne crosstalk noise in a data BWG of any PNoC occurs when the signal travels through the maximum number of intermediate nodes on its path to the WCPLN (because in this case the individual contributions of homodyne noise from all the intermediate nodes add up to incur the worst noise at the WCPLN).




Fig. 9: Worst-case SNR comparison of data and control waveguides of Corona, data waveguides of Flexishare, and data waveguides of Firefly with *HCTM* and their respective baseline configurations.

Fig. 9 summarizes the worst-case SNR results. It can be observed that SNR values for PNoCs with *HCTM* (brown bars) are very close to ideal SNR values (blue bars). This corroborates the excellent capabilities of our proposed *HCTM* technique in reducing homodyne crosstalk noise. From the figure, it can be observed that data BWGs of Corona, Flexishare and Firefly PNoCs with *HCTM* have 37.6%, 12.2%, and 4.6% SNR improvements on average compared to their respective baselines. Further, it can also be seen that the worst-case SNR of the Corona control BWG with *HCTM* increases by 20.1% compared to its baseline.




Fig. 10: Comparison of normalized energy consumption in Corona, Firefly and Flexishare PNoCs with *HCTM* and their respective baselines, for 12 PARSEC benchmarks. Energy consumption values are normalized to energy values of baseline configurations.

6.3 System-Level Overhead Analysis

In Section 5.2, we presented area and energy overhead of TDWGs. In this section, we present how the area and energy overheads of individual TDWGs manifest at the system-level. To know how the energy overhead of TDWGs affect the total energy consumption of the overall system, we perform simulation-based quantification of energy consumption for the aforementioned three PNoCs when used with *HCTM*. In our simulation study, we take 15dB as the target value of decoupled power for the TDWGs of the data BWGs of Firefly and Flexishare PNoCs, which is the same as chosen for the Corona data BWGs in Section 5.1. The results of this simulation study are plotted in Fig. 10. As evident from the figure, the Corona, Flexishare, and Firefly configurations with our *HCTM* technique have 19.2%, 18.5% and 0.5% higher energy consumption on average compared to their baseline configurations respectively. Total energy consumption for PNoCs with *HCTM* increases due to the additional dynamic energy and static energy consumption of the constituent TDWGs. As explained in section 5.2, the TDWGs corresponding to detector MRs consume dynamic energy every time a data packet is transferred over their corresponding BWG. The Firefly architecture, being a hybrid electro-photonic

NoC, routes fewer packets through photonic BWGs compared to Corona and Flexishare PNoCs, which in turn reduces the dynamic energy consumption in its TDWGs resulting in reduced total energy overhead.

Lastly, we took the physical dimensions of MRs and splitters from [32], and used them with the dimensions of TDWGs (from Section 5.2) to evaluate the photonic area overhead for *HCTM*. We found the photonic area overhead for *HCTM* in the Corona, Flexishare and Firefly PNoCs to be 6.76mm², 2.82mm², and 5.63mm² respectively. This area overhead is small relative to the total planar area of the chip (400 mm²).

7. CONCLUSIONS

In this paper, we demonstrated that the circulating power of MRs, built up due their resonant nature, induces time-dependent homodyne crosstalk noise in DWDM PNoCs. We also presented a lightweight low overhead homodyne crosstalk mitigation (*HCTM*) technique for the reduction of homodyne crosstalk noise in DWDM PNoCs. The *HCTM* technique is agnostic to the time-dependent characteristics of the homodyne crosstalk. Moreover, it also shows interesting trade-offs between reliability and energy overhead. We evaluate the effectiveness and overhead of our *HCTM* technique by implementing it for well-known PNoC architectures, including Corona, Firefly and Flexishare. Our experimental analysis shows that our approach when implemented on these PNoCs can improve the worst-case SNR by up to 37.6% compared to the baseline versions of these PNoCs, thereby significantly enhancing reliability, at the cost of up to 19.2% energy overhead and 1.7% photonic area overhead. Thus, *HCTM* represents an attractive solution to enhance reliability in emerging DWDM-based PNoCs.

ACKNOWLEDGMENTS

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

REFERENCES

- [1] D. Vantrease et al., "Corona: System Implications of Emerging Nanophotonic Technology," in Proc. *ISCA*, 2008.
- [2] Y. Pan et al., "Firefly: Illuminating Future Network-on-chip with Nanophotonics," in Proc. *ISCA*, 2009.
- [3] Y. Pan et al., "FlexiShare: Channel Sharing for an Energy-Efficient Nanophotonic Crossbar," in Proc. *HPCA*, 2010.
- [4] L. Duong et al., "Coherent Crosstalk Noise Analyses in Ring-Based Optical Interconnects," in Proc. *DATE*, 2015.
- [5] M. Nikdast et al., "Fat-Tree-Based Optical Interconnection Networks Under Crosstalk Noise Constraint," *IEEE TVLSI*, pp. 156–169, 2015.
- [6] Y. Xie et al., "Crosstalk Noise and Bit Error Rate Analysis for Optical Network-On-Chip," in Proc. *DAC*, 2010.
- [7] S. Chittamuru et al., "Crosstalk Mitigation for High-Radix and Low-Diameter Photonic NoC Architectures," *IEEE D&T*, pp. 29–39, 2015.
- [8] S. Dods et al., "Homodyne Crosstalk in WDM Ring and Bus Networks," *IEEE PTL*, pp. 1285–1287, 1997.
- [9] E. Tangdionga et al., "Experimental Evaluation of Optical Crosstalk Mitigation Using Phase Scrambling," *IEEE PTL*, pp. 567–569, 2000.
- [10] Y. Shen et al., "Coherent and Incoherent Crosstalk in WDM Optical Networks," *JLT*, pp. 759–764, 1999.
- [11] M. Bahadori et al., "Optimization of Microring-based Filters for Dense WDM Silicon Photonic Interconnects," *Measurement*, pp. 5.
- [12] F. Xia et al., "Ultra-compact high order ring resonator filters using submicron silicon photonic wires for on-chip optical interconnects," *Opt. Express*, pp. 11934, 2007.
- [13] Q. Li et al., "Design and Demonstration of Compact, Wide Bandwidth Coupled-Resonator Filters on a Silicon-On-Insulator Platform," *Opt. Express*, pp. 2247, 2009.

- [14] M. Nikdast et al., "Systematic Analysis of Crosstalk Noise in Folded-Torus-Based Optical Networks-on-Chip," *IEEE TCAD*, 2014.
- [15] Y. Xie et al., "Formal Worst-Case Analysis of Crosstalk Noise in Mesh-Based Optical Networks-on-Chip," *IEEE TVLSI*, pp. 1823–1836, 2013.
- [16] B. Lin et al., "Crosstalk Analysis for Microring Based Optical Interconnection Networks," *JLT*, pp. 2415–2420, 2012.
- [17] W. Bogaerts et al., "Silicon Microring Resonators," *Laser Phot.*, 2012.
- [18] *Optical Microresonators*, vol. 138. Springer New York, 2008.
- [19] R. Wu et al., "Compact Models for Carrier-Injection Silicon Microring Modulators," *Opt. Express*, pp. 15545, 2015.
- [20] K. Vahala et al., "Optical microcavities," *Nature*, pp. 839–846, 2003.
- [21] R. Soref et al., "Electrooptical Effects in Silicon," *IEEE JSTQE*, pp. 123–129, 1987.
- [22] K. Padmaraju et al., "Resolving the Thermal Challenges for Silicon Microring Resonator Devices," *Nanophotonics*, pp. 269–281, 2014.
- [23] Q. Xu et al., "Silicon Microring Resonators with 1.5- μm Radius," *Opt. Express*, pp. 4309, 2008.
- [24] Y. Boucher et al., "Analytical Model for the Coupling Constant of a Directional Coupler in Terms of Slab Waveguides," *Opt. Eng.*, 2014.
- [25] C. Bienia et al., "The PARSEC Benchmark Suite: Characterization and Architectural Implications," in *Proc. PACT*, 2008.
- [26] N. Binkert et al., "The Gem5 Simulator," *Comput Arch. News*, 2011.
- [27] C. Sun et al., "DSENT - A Tool Connecting Emerging Photonics with Electronics for Opto-Electronic Networks-On-Chip Modeling," in *NOCS*, 2012.
- [28] L. Duong et al., "A Case Study of Signal-to-Noise Ratio in Ring-Based Optical Networks-on-Chip," *IEEE MDAT*, pp. 55–65, 2014.
- [29] X. Zheng et al., "Ultra-Efficient 10Gb/s Hybrid Integrated Silicon Photonic Transmitter and Receiver," *Opt. Express*, pp. 5172, 2011.
- [30] P. Grani et al., "Design Options for Optical Ring Interconnect in Future Client Devices," *JETC* 2014.
- [31] C. Nitta et al., "Addressing System-Level Trimming Issues in On-Chip Nanophotonic Networks," in *Proc. HPCA*, 2011.
- [32] S. Xiao et al., "Modeling and measurement of losses in silicon-on-insulator resonators and bends," *Opt. Express*, pp. 10553, 2007.
- [33] N. Sherwood-Droz et al., "Scalable 3D Dense Integration of Photonics on Bulk Silicon," in *Opt. Express*, p. 17758, 2011.
- [34] S. Chittamuru, I. Thakkar, S. Pasricha, "PICO: Mitigating Heterodyne Crosstalk Due to Process Variations and Intermodulation Effects in Photonic NoCs", in *DAC*, 2016.
- [35] S. Chittamuru, I. Thakkar, S. Pasricha, "Process Variation Aware Crosstalk Mitigation for DWDM based Photonic NoC Architectures", in *ISQED*, 2016.
- [36] I. Thakkar, S. Chittamuru, S. Pasricha, "A Comparative Analysis of Front-End and Back-End Compatible Silicon Photonic On-Chip Interconnects", in *SLIP*, 2016.
- [37] S. Bahirat, S. Pasricha, "Exploring Hybrid Photonic Networks-on-Chip for Emerging Chip Multiprocessors", *IEEE/ACM CODES+ISSS*, 2009.
- [38] S. Pasricha, S. Bahirat, "OPAL: A Multi-Layer Hybrid Photonic NoC for 3D ICs", *IEEE/ACM ASPDAC*, 2011.
- [39] S. Pasricha, N. Dutt, "Trends in Emerging On-Chip Interconnect Technologies", *IPSI Transactions on System LSI Design Methodology*, Vol. 1, Sep 2008.