

HYDRA: Heterodyne Crosstalk Mitigation With Double Microring Resonators and Data Encoding for Photonic NoCs

Sai Vineel Reddy Chittamuru, *Student Member, IEEE*, Ishan G. Thakkar, *Student Member, IEEE*,
and Sudeep Pasricha, *Senior Member, IEEE*

Abstract—Silicon-photonic networks on chip (PNoCs) provide high bandwidth with lower data-dependent power dissipation than does the traditional electrical NoCs (ENoCs); therefore, they are promising candidates to replace ENoCs in future manycore chips. PNoCs typically employ photonic waveguides with dense wavelength division multiplexing (DWDM) for signal traversal and microring resonators (MRs) for signal modulation. Unfortunately, DWDM increases susceptibility to intermodulation (IM) and off-resonance filtering effects, which reduce optical signal-to-noise ratio (OSNR) for photonic data transfers. Additionally, process variations (PVs) induce variations in the width and thickness of MRs causing resonance wavelength shifts, which further reduce OSNR, and create communication errors. This paper proposes a novel cross-layer framework called HYDRA to mitigate heterodyne crosstalk due to PVs, off-resonance filtering, and IM effects in PNoCs. The framework consists of two device-level mechanisms and a circuit-level mechanism to improve heterodyne crosstalk resilience in PNoCs. Simulation results on three PNoC architectures indicate that HYDRA can improve the worst case OSNR by up to 5.3 \times and significantly enhance the reliability of DWDM-based PNoC architectures.

Index Terms—Crosstalk noise, photonic NoCs (PNoCs), process variations (PVs).

I. INTRODUCTION

WITH several hundreds of on-chip cores expected to become a reality in the near future, electrical networks on chip (ENoCs) are projected to suffer from very high-power dissipation and poor performance. Recent developments in the area of 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 (see [1]–[5]). These architectures employ on-chip photonic links that use microring resonator (MR) modulators to modulate electrical signals onto photonic signals that travel through a silicon waveguide, and receivers that

are comprised of MR filters and photodetectors coupled on their drop ports. The MR filters of these receivers filter and drop photonic signals on to the photodetectors to recover electrical signals. Each MR has a unique set of (resonance) wavelengths that it can couple to and work correctly with. Dense wavelength division multiplexing (DWDM) enables multiple single-wavelength communication channels to share a silicon waveguide, to achieve high bandwidth parallel on-chip data transfers. The periodic passband of MRs that are in resonance with these wavelengths constrains the degree of DWDM (i.e., number of DWDM wavelength channels per waveguide). Unfortunately, MRs suffer from intrinsic crosstalk noise and power loss due to their design imperfections.

Duong *et al.* [6] categorizes crosstalk noise into two types: homodyne (coherent) and heterodyne (incoherent). The homodyne crosstalk noise power of a particular wavelength affects the signal power of the same wavelength, whereas with heterodyne crosstalk the signal power gets affected by some noise power of one or more other (different) wavelengths. Heterodyne crosstalk is a major contributor of noise in DWDM-based PNoCs, and reduces optical signal-to-noise ratio (OSNR) and reliability in PNoCs [6].

Due to the heterodyne crosstalk phenomenon, when a data-modulated wavelength passes by an MR, depending on its data bitrate (modulation rate), average spectral power, and its relative detuning from the resonance of the MR, a part of its power is dropped by the MR [41]. All modulator, filter, and switch MRs can drop signal power due to heterodyne crosstalk. This heterodyne crosstalk-induced signal power drop creates impairments in the passing nonresonant signals. These impairments in a signal result in smoothed transition edges, lengthened rise and fall times, dampened signal amplitude, suppressed signal strength, and reduced extinction ratio, which cause data errors in the signal [44]. The overall impact of these signal impairments is manifested as a power penalty, which is defined as the amount of extra power required at the detector to overcome the data errors caused by these signal impairments.

Heterodyne crosstalk-induced signal power drop has an additional effect, referred to as off-resonance filtering, at the filter MRs that are coupled with detectors. When a filter MR drops some power from the adjacent nonresonant signals onto a detector at its drop port, this dropped optical power (i.e., crosstalk noise power) produces proportional (pessimistic case) or shot-noise limited (optimistic case) noise current in the detector. This noise current increases the noise floor of the

Manuscript received January 25, 2017; revised May 7, 2017 and July 18, 2017; accepted August 21, 2017. Date of publication September 18, 2017; date of current version December 27, 2017. This work was supported in part by SRC, in part by the U.S. National Science Foundation under Grant CCF-1252500 and Grant CCF-1302693, and in part by AFOSR under Grant FA9550-13-1-0110. (Sai Vineel Reddy Chittamuru and Ishan G. Thakkar contributed equally to this work.) (Corresponding author: Sai Vineel Reddy Chittamuru.)

The authors are with the Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO 80523 USA (e-mail: sai.chittamuru@colostate.edu; ishan.thakkar@colostate.edu; sudeep@colostate.edu).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/TVLSI.2017.2749967

detector, increasing the minimum detectable signal power for the detector. As a result, the detector requires larger signal power to achieve a target OSNR in the presence of this crosstalk noise power. One of our goals is to reduce crosstalk noise power in detectors due to this off-resonance filtering effect.

The strength of the heterodyne crosstalk noise power at a detector depends on the following three attributes: 1) channel gap (CG) between the MR resonant wavelength and the adjacent wavelength signals; 2) Q-factors of neighboring detector-coupled filter MRs, and 3) the strengths of the nonresonant signals at the detector-coupled filter MR. With increase in DWDM, the CG between two adjacent wavelength signals decreases, which in turn increases heterodyne crosstalk noise power in detectors. With decrease in Quality factors (Q-factors) of MRs, the widths of the resonant passbands of MRs increase, increasing passband overlap with neighboring nonresonant signals, which in turn increases heterodyne crosstalk noise power. The strengths of the nonresonant signals depend on the losses faced by the nonresonant signals throughout their path from the laser source to the detector-coupled MR filter.

Intermodulation (IM) crosstalk has the biggest influence on the last attribute discussed above, causing suppression (or loss) of signal strength of nonresonant signals in a DWDM waveguide [7]. IM crosstalk occurs when a modulator MR induces impairments in, and as a result, suppresses the neighboring nonresonant signals. Thus the level of heterodyne crosstalk noise power and resultant OSNR at the detector depends on the amount of IM crosstalk-induced signal suppression at the modulator. This motivates mitigating the effects of IM crosstalk-induced signal suppression on heterodyne crosstalk by controlling the strengths of the nonresonant signals at the detector.

Additionally, fabrication process variations (PVs) induce variations in the width, thickness, and doping concentration width and thickness of active MRs, which cause resonance wavelength shifts in MRs [8]–[9]. PV-induced resonance shifts, when uncompensated, may reduce the gap between the resonances of the victim MRs and adjacent MRs, which increase crosstalk and worsen OSNR. For example, a previous study shows that in a DWDM-based photonic link with 1.48-nm channel spacing and 4-Gbps bitrate, when PV-induced resonance shift is over 1/3rd of the CG, bit error rate increases from 10^{-12} to 10^{-6} [10]. Techniques to counteract PV-induced resonance shifts in MRs involve realigning the resonant wavelengths by using localized trimming [11] or thermal tuning [12]. Localized trimming induces a blue shift in the resonance wavelengths (to compensate PV-induced red shifts) of MRs using carrier injection into MRs, whereas thermal tuning induces a red shift in the resonance wavelengths (to compensate PV-induced blue shifts) of MRs through heating or thermal tuning of MRs using microheaters. However, our analysis has shown that localized trimming and thermal tuning increase intrinsic optical loss in MRs and signal loss in waveguides due to the free carrier absorption effect (FCA) [24] and increased optical scattering [36]. It is important to address this increase in loss, which drives the MR away from critical

coupling and decreases its Q-factor, increasing heterodyne crosstalk and reducing OSNR [35].

In this paper, we present a novel cross-layer heterodyne crosstalk-mitigation framework called HYDRA to address the abovementioned challenges and enable reliable communication in emerging PNoC-based manycore chips. Our framework has low overhead and is easily implementable on any existing DWDM-based PNoC without major modifications to the architecture. Our novel contributions are as follows.

- 1) We present device-level analytical models to capture the deleterious effects of localized trimming and thermal tuning in MRs. We also extend these models for system-level heterodyne crosstalk analysis.
- 2) We propose a device-level method for IM effect-induced signal suppression aware heterodyne crosstalk mitigation (IMCM) that improves worst case OSNR in detectors by controlling the nonresonant signal power.
- 3) We propose another device-level technique for heterodyne crosstalk mitigation [double-MR-based crosstalk mitigation (DMCM)] that uses double MRs (DMRs) to improve worst case OSNR in detectors by tailoring the MRs' passbands to have steeper roll-off.
- 4) We propose a circuit-level technique for heterodyne crosstalk mitigation (EDCM) that improves worst case OSNR in detectors by encoding data to avoid undesirable data value occurrences.
- 5) We combine IMCM, DMCM, and EDCM into a holistic cross-layer heterodyne crosstalk-mitigation framework called HYDRA and evaluate it on three well-known crossbar PNoC architectures as well as prior work on heterodyne crosstalk mitigation.

II. BACKGROUND AND RELATED WORK

DWDM-based PNoCs [49]–[51] utilize several photonic devices such as MRs, waveguides along which data is transmitted on multiple wavelengths using DWDM, topology-specific photonic path splitters, photodetectors to detect photonic wavelength signals and convert them into proportional electrical signals, and trans-impedance amplifiers for detected signal amplification. Each MR has a unique resonance wavelength in the utilized DWDM spectrum that it can couple to and work correctly with. This resonant nature of an MR allows it to be used as a filter or a switch. A filter MR is used to filter and drop its resonance wavelength on to a photodetector, whereas a switch MR is used to route the propagation of a resonant wavelength signal between two waveguides. Typically, an MR can electro-optically be driven ON and OFF resonance with its resonance wavelength, which allows the MR to modulate 1s (when OFF resonance) and 0s (when ON-resonance) on its resonance wavelength. The reader is referred to [18], [21], and [22] for more details on these devices.

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 losses such as insertion losses in MR modulators and filters [43], propagation and bending loss in waveguides, and splitting loss in splitters. Such losses negatively impact signal strength in waveguides, which

reduces OSNR for a given noise power. In addition to the optical signal loss, crosstalk noise of the constituent MRs also deteriorates OSNR. Crosstalk noise in PNoCs usually occurs due to the imperfections in MRs used as optical modulators, filters, and switches. This crosstalk noise can be classified as homodyne or heterodyne.

For homodyne crosstalk, the noise power has the same wavelength as the signal power. As demonstrated in [4], out-of-phase homodyne crosstalk noise always degrades signal integrity. Homodyne crosstalk may either contribute to noise or cause fluctuations in signal power, which makes the analysis and mitigation of homodyne crosstalk complicated and beyond the scope of this paper. On the other hand, heterodyne crosstalk occurs when an MR picks up some optical power from nonresonant signals (as explained in Section I). This paper proposes solutions to mitigate heterodyne crosstalk due to the off-resonance filtering effect. In the rest of this paper, we use the term crosstalk to refer to heterodyne crosstalk, unless specified otherwise.

A few prior works have analyzed crosstalk in PNoCs. The effect of crosstalk noise on OSNR is shown to be negligible in the WDM system presented in [14], as this system uses only four WDM wavelengths per waveguide with 1.3-nm channel spacing and 4-Gbps bitrate. In [7], IM crosstalk is shown to be negligible for a WDM link operating at 10 Gbps with a channel spacing of 1.6 nm. However, in PNoC architectures that use DWDM (e.g., Corona [1] with 64-wavelength DWDM), significant crosstalk noise is expected. The damaging impact of crosstalk noise in the Corona PNoC is presented in [16], where worst case OSNR is estimated to be 14 dB in data waveguides, which is insufficient for reliable data transfers. To mitigate the impact of crosstalk noise in DWDM-based PNoC architectures, two encoding techniques and one wavelength spacing technique were presented in [13] and [17]. However, none of these works considers the system-level impact of IM effects, OFF-resonance filtering, or PVs on crosstalk noise in DWDM-based PNoCs.

Fabrication-induced PVs impact the cross section, i.e., width and height, of photonic devices such as MRs and waveguides. In MRs, PV causes resonance wavelength drifts, which can be counteracted by using device-level techniques such as localized trimming [11] and thermal tuning [12]. Trimming induces a blue shift in the resonance wavelengths of MRs using carrier injection into MRs, whereas thermal tuning induces a red shift in the resonance wavelengths of MRs through heating of MRs using ring heaters. Such device-level techniques are essential to overcome PV-induced drifts, but they incur high-power overheads and may increase signal loss and crosstalk noise, thereby reducing OSNR. This motivates the use of supplementary system-level approaches to reduce the overheads of device-level techniques. A few prior works have explored the impact of PV on DWDM-based PNoCs at the system level [10], [20]. In [10], a thermal tuning-based approach is presented that adjusts chip temperature using dynamic voltage and frequency scaling to compensate for chip-wide PV-induced resonance shifts in MRs. In [20], a methodology to salvage network-bandwidth loss due to PV-drifts is proposed, which reorders MRs and trims them

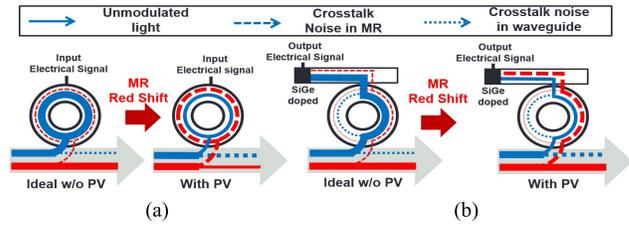


Fig. 1. Impact of PV-induced resonance shifts on MR operation in DWDM waveguides (note: only PV-induced red resonance shifts are shown). (a) MR as active modulator with PV-induced red shift, modulating in resonance wavelength. (b) Detector-coupled MR filter with PV-induced red shift, filtering its resonance wavelength and dropping it on the detector.

to nearby wavelengths. But the achievable benefits for all these supplementary system-level techniques highly depend on the underlying system architecture and they also ignore the harmful effects of device-level PV remedies (i.e., trimming and tuning) on crosstalk.

III. PV-AWARE CROSSTALK ANALYSIS

An MR can be considered to be a looped photonic waveguide with a small diameter, not to be confused with the straight photonic waveguide used for wavelength-parallel data transfers for which MRs serve as modulators and filters. Variations in an MR's dimensions due to PV cause a "shift" in its resonance wavelength. Fig. 1 shows the impact of PV on crosstalk noise (dashed lines) in MRs. From Fig. 1(a), PV-induced red shifts in MR modulators increase crosstalk noise in the waveguide and decrease signal strength of nonresonating wavelength signals. Fig. 1(b) shows how PV-induced red shifts increase detected crosstalk noise and decrease detected signal power of resonance wavelengths in detectors, which in turn reduces OSNR and photonic data communication reliability. As discussed earlier, localized trimming and thermal tuning are essential to deal with PV-induced resonance red and blue shifts in MRs, respectively. However, the use of these methods in an MR alters its intrinsic optical properties, which leads to increased crosstalk and degraded performance in PNoCs that use these MRs.

A. Impact of Localized Trimming on Crosstalk

The localized trimming method injects extra free carriers in the circular MR waveguide to counteract the PV-induced resonance red shift. The introduction of extra free carriers reduces the refractive index of the looped MR waveguide, which in turn induces a blue shift in resonance to counteract the PV-induced red shift. However, the extra free carriers increase the absorption related optical loss in the MR due to the FCA [24]. The increase in optical loss results in a decrease of MR Q-factor, which increases MR insertion loss and crosstalk. We use a PV map (described in more detail in Section IV) 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 (Δn_{Si}) required to counteract this PV-induced

wavelength shift using the following equation [27]:

$$\Delta\lambda_r = \frac{\Delta n_{\text{eff}} * \lambda_r}{n_g} \approx \frac{\Gamma * \Delta n_{\text{Si}} * \lambda_r}{n_g} \quad (1)$$

where $\Delta\lambda_r$ is the PV-induced resonance shift that needs to be compensated for, λ_r is the target resonance wavelength of the MR, and n_g is the group refractive index (ratio of speed of light to group velocity of all wavelengths traversing the waveguide) of the MR waveguide. Moreover, Δn_{eff} is the change in effective index that is approximately equal to $\Gamma \times \Delta n_{\text{Si}}$ [39], where Γ is the confinement factor describing the overlap of the optical mode with the MR waveguide's silicon core. The waveguides used in this paper (both MRs' looped waveguides and straight bus waveguides) are rectangular channel waveguides fabricated using Si-SiO₂ material with a cross section of 450 nm × 220 nm. We model these waveguides using a commercial eigenmode solver [40], based on which the values of Γ and n_g at 1550 nm are calculated to be 0.78 and 4.16, respectively.

The change in free carrier concentration required to induce refractive index change of Δn_{Si} at around 1.55- μm wavelength can be quantified as follows [24]:

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

where ΔN_e and ΔN_h are the change in free electron concentration and free hole concentration, respectively. The change in the absorption loss coefficient ($\Delta\alpha_{\text{Si}}$) due to the change in free carrier concentration (owing to the FCA effect) can be quantified using the following equation [24]:

$$\Delta\alpha_{\text{Si}} = 8.5 \times 10^{-18} \Delta N_e + 6.0 \times 10^{-18} \Delta N_h. \quad (3)$$

Q-factor is a measure of the sharpness of the MR's resonance relative to its central (resonant) wavelength [27]. The Q-factor of MRs affects the magnitudes of crosstalk penalties (as explained in [7] and [41]) and determines the photon-lifetime limited allowable bitrate of signals [42]. Moreover, the Q-factor of an MR represents the number of oscillations of the field in the MR before the circulating field energy in the MR is depleted to 1/e of the initial energy [27]. Now, from [27], the field-energy decay in the MR cavity depends on the losses in the cavity. Therefore, the Q-factor of an MR depends on the MR's loss coefficient (α) along with some other parameters. The relationship between the Q-factor and the change in absorption loss coefficient ($\Delta\alpha_{\text{Si}}$) is given by the following equations [27]:

$$Q' = Q + \Delta Q = \frac{2\pi^2 R n_g \sqrt{r_1 r_2 a'}}{\lambda_r (1 - r_1 r_2 a')} \quad (4)$$

$$a' = a + \Delta a = e^{-\pi R(\alpha + \Gamma \Delta\alpha_{\text{Si}})} \quad (5)$$

where r_1 and r_2 are the self-coupling coefficients of an add-drop MR (defined in [27]); R is the MR radius; a' is the resultant round-trip field transmission after an arbitrary change Δa in the original round-trip field transmission a ; $\Delta\alpha_{\text{Si}}$ is the change in the MR's original loss coefficient α ; and ΔQ is the change in the loaded Q-factor (Q). Equation (4) gives the resultant loaded Q-factor Q' for an add-drop MR. Similarly, the Q' for an all-pass MR (described in [27]) can be modeled

by setting $r_2 = 1$ in (4). Note that, as depicted in Fig. 1, we use all-pass MRs as modulators and add-drop MRs as filters and switches.

Now, the original loss coefficient α is a sum of three components: 1) intrinsic loss coefficient due to material loss and sidewall roughness-induced scattering loss; 2) bending loss coefficient, which is a result of the curvature in the MR; and 3) the absorption effect factor that depends on the original free carrier concentration in the waveguide core. Typically, the localized trimming method (when used to induce a blue shift in the MR resonance) injects excess concentration of free carriers into the MR, which increases the absorption loss coefficient (positive $\Delta\alpha_{\text{Si}}$). As evident from (5), a positive value of $\Delta\alpha_{\text{Si}}$ results in a decrease in a' , which in turn decreases the Q-factor Q' [from (4)]. This causes a broadening of the MR passband, which results in increased insertion loss, crosstalk noise, and signal impairment/degradation related power penalty.

We model the MR transmission spectrum using a Lorentzian function [26]. In (6), this function is used to represent coupling factor ϕ [6] between wavelength λ_i and an MR with resonance wavelength λ_j . From [6], we use this coupling factor ϕ to model the heterodyne crosstalk noise power (of wavelength λ_i) that is dropped on the detector at the drop port of a filter MR with resonance wavelength λ_j . From [7], IM crosstalk incurred by a modulator MR induces signal impairment, suppressing the power in the adjacent signal. As in [7], we use the same Lorentzian function to determine a loss factor γ in (7), which is the factor by which signal power of a wavelength λ_i is suppressed when it passes by a modulator MR whose resonance wavelength is λ_j . Thus, when a wavelength signal in a waveguide passes by a modulator MR, the IM-crosstalk-induced bitrate independent suppression in its power can be modeled as a through loss, which is defined as γ times the signal power before it passes by the MR.

Now from (3)–(5), Q' of an MR decreases with localized trimming-based increase in carrier concentration. This in turn increases ϕ and crosstalk noise power (6). Note that we do not consider the effect of decrease in free carrier concentration, as we use only carrier injection for both modulation and trimming (to counteract PV-induced red shifts). As would be clear in Section III-B, we do not need to use carrier depletion with trimming, as we would rather heat up the MRs at higher temperatures to counter the PV-induced blue shifts

$$\Phi(\lambda_i, \lambda_j, Q') = \left(1 + \left(\frac{2Q'(\lambda_i - \lambda_j)}{\lambda_j} \right)^2 \right)^{-1} \quad (6)$$

$$\gamma(\lambda_i, \lambda_j, Q') = \left(1 + \left(\frac{2Q'(\lambda_i - \lambda_j)}{\lambda_j} \right)^{-2} \right)^{-1}. \quad (7)$$

B. Impact of Thermal Tuning of MR on Crosstalk

As mentioned earlier, localized trimming-based carrier injection induces blue shifts in resonance wavelengths of MRs, which can be used to compensate PV-induced red shifts in resonance wavelengths. In contrast, thermal tuning of MRs incurs red shifts in resonance wavelengths of MRs, which can

be used to compensate PV-induced blue shifts in resonance wavelengths. From Section III-A, localized trimming results in increased absorption loss coefficient, subsequent decrease in Q-factor, and increase in insertion loss and crosstalk power penalties. Similarly, it can be intuitively inferred that heating of MRs would also increase the absorption loss coefficient in MRs, because, the increase in temperature from the heating of MRs imparts enough energy to some valence electrons of doped silicon (constitutive semiconductor material of MRs) so that they become free carriers. However, these extra free electrons do not significantly increase the net concentration of free carriers in doped silicon. This is because, in doped silicon, the majority of free carriers emanate from the ionization of dopant atoms and usually, all the dopant atoms are completely ionized at room temperature [19]. Thus, any increase in the MR operating temperature above room temperature does not cause ionization of any more dopant atoms. As a result, the concentration of the majority free carriers, and hence, the net free carrier concentration in doped silicon does not change with heating of MRs. Therefore, heating of MRs does not increase the absorption loss coefficient of MRs.

The scattering loss coefficient (that gives fractional loss in signal amplitude) of an MR's circular waveguide is proportional to the refractive index contrast between the core and the cladding ($n_{Si} - n_{SiO_2}$) of the MR waveguide and the size of the surface roughness σ , and is given by the following equation [36], [37]:

$$\alpha_{scatter} = \frac{4(\cos\theta)^3 k_0^2 n_1^2 \sigma^2}{\sin\theta} \cdot \left(\frac{k_0 \sqrt{n_1^2 (\sin\theta)^2 - n_2^2}}{L k_0 \sqrt{n_1^2 (\sin\theta)^2 - n_2^2} + 2} \right) \quad (8)$$

where $\alpha_{scatter}$ is scattering loss coefficient, k_0 is the free-space wavenumber, $n_1 = n_{Si}$ is the MR core's refractive index, $n_2 = n_{SiO_2}$ is the MR cladding's refractive index, L is the MR thickness, and θ is the propagation angle for the fundamental mode in the MR. With heating of the MR, the refractive index n_{Si} (of the MR's core) and the refractive index n_{SiO_2} (of the MR's cladding) increase to their new values of $n_{Si} + \Delta n_{Si}$ and $n_{SiO_2} + \Delta n_{SiO_2}$, respectively, which are given by the following equations:

$$n_{Si}^{T+\Delta T} = n_{Si}^T + \Delta n_{Si} = n_{Si}^T + \frac{\delta n_{Si}}{\delta T} \cdot \Delta T \quad (9)$$

$$n_{SiO_2}^{T+\Delta T} = n_{SiO_2}^T + \frac{\delta n_{SiO_2}}{\delta T} \cdot \Delta T \quad (10)$$

where $\delta n_{Si}/\delta T$ and $\delta n_{SiO_2}/\delta T$ are the thermo-optic coefficients of Si (MR's core) and SiO_2 (MR's cladding) materials, respectively, and they assume the values of $1.86 \times 10^{-4} \text{ K}^{-1}$ and $1 \times 10^{-5} \text{ K}^{-1}$, respectively [34]. ΔT is an increase in temperature of the MR due to heating. Due to smaller thermo-optic coefficient of SiO_2 and smaller mode field confinement in SiO_2 cladding, the effects of temperature change on $n_{SiO_2}^{T+\Delta T}$ is negligible. If a blue shift in an MR's resonance wavelength of $\Delta\lambda_R$ is to be compensated by heating the MR, the required increase in MR's temperature can be computed using the following equation [34]:

$$\Delta\lambda_r = \Gamma \cdot \frac{\delta n_{Si}}{\delta T} \cdot \frac{\lambda_r}{n_g} \cdot \Delta T. \quad (11)$$

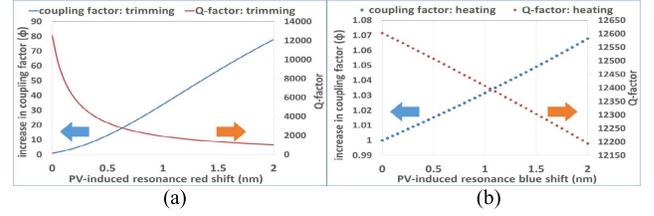


Fig. 2. (a) Effect of localized trimming. (b) Effect of thermal tuning, on the Q-factor and fractional increase in coupling factor of an example MR. Here, the fractional increase in coupling factor is calculated with respect to the original coupling factor of the MR without PV.

Now, as the thermo-optic coefficient of Si is greater than that of SiO_2 , $n_{Si}^{T+\Delta T}$ increases faster with increase in temperature than $n_{SiO_2}^{T+\Delta T}$. As a result, the difference $(n_1^2 (\sin\theta)^2 - n_2^2)$ in (8), which depends on the index contrast between the core and the cladding, increases with increase in temperature. This leads to an increase in $\alpha_{scatter}$ with increase in temperature (8). Now, similar to the case of localized trimming, this increase in scattering loss coefficient leads to decrease in MR Q-factor. Using (8), the increased value of scattering loss coefficient $\alpha_{scatter}$ can be calculated, which then can be used in place of $(\alpha + \Delta\alpha_{Si})$ in (5) to find the decreased value of Q-factor from (4).

To model and compare the effects of localized trimming and thermal tuning of MRs, we simulate an MR with a radius (R) of $1.8 \mu\text{m}$ (deemed as implementable with CMOS-type processes based on projections from [23]) considering initial original Q-factor of 12500, self-coupling coefficients $r_1 = 0.99$, $r_2 = 0.99$, and field-transmission coefficient a of 0.991. Note that we use the initial $Q = 12500$, because it gives the optimum value of total MR filter penalty for 5-Gbps bitrate and 64 channels (as projected from [41]). Also, note that we assume initial $\alpha_{scatter} = 0.14 \text{ cm}^{-1}$, which corresponds to $\sigma = 1 \text{ nm}$, $n_{Si} = 3.5$, $n_{SiO_2} = 1.5$, $L = 220 \text{ nm}$, and $\theta = 26.51$ in (8).

Using (1)–(10), we evaluate the values of Q-factor and increase in coupling factor ϕ for this example MR, when PV-induced red/blue shifts of different values in the resonance wavelength of this MR are compensated by using localized trimming/thermal tuning. Fig. 2(a) and (b) plots these values of Q-factor and ϕ for localized trimming and thermal tuning, respectively. From Fig. 2, compensating 2-nm PV-induced red shift in an MR's resonance wavelength with localized trimming decreases the MR's Q-factor by 91.7% and increases ϕ by $77.8\times$ compared to original Q-factor and coupling factor, respectively. Furthermore, compensating 2 nm of PV-induced blue shift in MR's resonance wavelength with thermal tuning decreases the MR's Q-factor by only 3.25% and increases ϕ by $1.07\times$ compared to the original Q-factor and coupling factor, respectively. Thus, it can be concluded that thermal tuning of MRs has negligible impact on MRs' Q-factor and coupling factor compared to localized trimming. Therefore, compared to localized trimming, thermal tuning does not significantly increase insertion loss and crosstalk penalties for MRs.

However, note that thermal tuning cannot compensate for PV-induced red shifts in MRs' resonance wavelengths. Therefore, in a typical PNoC, where both red and blue shifts in

TABLE I

PHOTONIC POWER LOSS AND CROSSTALK COEFFICIENTS [16], [33], [45]

Notation	Parameter type	Parameter value (in dB)
L_P	Propagation loss	-0.274 per cm
L_B	Bending loss	-0.0085 per 90°
L_{S12}	1×2 splitter power loss	-0.2
L_{S14}	1×4 splitter power loss	-0.2
L_{S16}	1×6 splitter power loss	-0.2

TABLE II

OTHER MODEL PARAMETER NOTATIONS [33]

Notation	Crosstalk Coefficient	Parameter Value
Q	Q-factor	9000
RS	Detector responsivity	0.8 A/W
L	Photonic path length in cm	
B	Number of bends in photonic path	
λ_r	Resonance wavelength of MR	
R_{S12}	Splitting factor for 1×2 splitter	
R_{S14}	Splitting factor for 1×4 splitter	
R_{S16}	Splitting factor for 1×6 splitter	

MRs' resonance wavelengths are present, the use of localized trimming is inevitable. As a result, it is imperative to overcome the poor efficiency of localized trimming. We propose, as part of our HYDRA framework, a circuit-level technique for heterodyne crosstalk mitigation (EDCM; Section VII) that mitigates the effect of PV-remedial techniques (both localized trimming and thermal tuning) on MR crosstalk penalties. Furthermore, this paper only analyzes the impact of PV and its remedial techniques on crosstalk noise. Evaluating the impact of thermal variations on crosstalk noise is beyond the scope of this paper. In the following section, we use the derived values of Φ and γ from this and the previous section to model worst case crosstalk and OSNR for the Corona PNoC, in the presence of PVs.

C. PV-Aware Crosstalk Models for Corona PNoC

We characterize crosstalk in waveguides with DWDM for the Corona PNoC enhanced with token-slot arbitration [1]. We present equations to model the OFF-resonance filtering effect-induced crosstalk noise power and resultant OSNR in the detectors of receiver groups. Before presenting actual equations, we show notations for parameters used in the 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 four cores in each cluster. A photonic crossbar topology with 64 data waveguide groups is used for communication between clusters. Each data waveguide group consists of four 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 four MWSR waveguides in a single cycle by a sender. Each of the 64 data waveguide groups starts at a different cluster called "home-cluster," traverses other clusters (where modulators can modulate light and receivers can filter and detect this light), and finally ends at the home-cluster again, at a set of receivers(optical termination).

A power waveguide supplies optical power from an off-chip laser to each of the 64 data waveguide groups at its home-cluster via a series of 1×2 splitters. In each of the 64 home-clusters, optical power is distributed among four MWSR waveguides equally using a 1×4 splitter with splitting factor R_{S14} . As all 1×2 splitters are present before the last (64th) waveguide group, this waveguide group suffers the highest signal power loss. Therefore, the worst case signal and crosstalk noise exists in the detectors of the receiver group of the 64th cluster node, and this node is called the worst case power loss node (N_{WCPL}) in the Corona PNoC.

For this N_{WCPL} node of the Corona PNoC, the signal power ($P_{\text{signal}}(\lambda_j)$) and crosstalk noise power ($P_{\text{noise}}(\lambda_j)$) received at a receiver (i.e., detector-coupled MR filter) with resonance wavelength λ_j are expressed in (12) and (13), respectively. $K(\lambda_i)$ in (14) represents signal power loss of λ_i before the receiver group of N_{WCPL} . $\psi(\lambda_i, \lambda_j)$ in (15) represents signal power loss of λ_i before the receiver with resonance wavelength λ_j within the receiver group of N_{WCPL} . $P_S(\lambda_i, \lambda_j)$ in (16) is the signal power of the λ_i wavelength in the waveguide that has reached the receiver with λ_j resonance wavelength in the receiver group of N_{WCPL} after passing through all the preceding receivers. Due to PV (more details about modeling of PV in PNoCs are presented in the following section), crosstalk coupling factor [ϕ , (6)] increases with decrease in loaded Q-factor [Q' , which is calculated by using (4) and (5)], which in turn increases OFF-resonance filtering effect induced crosstalk noise in the detectors. Furthermore, $Q'_{(x \times y)+j}$ is defined as the Q-factor of j th MR which is in the $x+1$ th node and each node is having "y" number of MRs. We can define OSNR(λ_j) at the detector in the receiver (with resonance wavelength λ_j) of N_{WCPL} as the ratio of $P_{\text{signal}}(\lambda_j)$ to $P_{\text{noise}}(\lambda_j)$, as shown in (17). These equations [i.e., (12)–(17)] are based on the models presented in [6] and [16]

$$P_{\text{signal}}(\lambda_j) = \Phi(\lambda_j, \lambda_j, Q'_{(63 \times 64)+j}) P_S(\lambda_j, \lambda_j) \quad (12)$$

$$P_{\text{noise}}(\lambda_j) = \sum_{i=1}^n \Phi(\lambda_i, \lambda_j, Q'_{(63 \times 64)+j}) (P_S(\lambda_i, \lambda_j)) (i \neq j) \quad (13)$$

$$K(\lambda_i) = (R_{S14})(L_{S14})(L_P)^L (L_B)^B \times \prod_{n=1}^{63} \prod_{j=1}^{64} \gamma(\lambda_i, \lambda_j, Q'_{((n-1) \times 64)+j}) \quad (14)$$

$$\psi(\lambda_i, \lambda_j) = \prod_{k=1}^{(k-1) < j} \gamma(\lambda_i, \lambda_k, Q'_{(63 \times 64)+k}) \quad (15)$$

$$P_S(\lambda_i, \lambda_j) = K(\lambda_i) \psi(\lambda_i, \lambda_j) P_{\text{in}}(i) \quad (16)$$

$$\text{OSNR}(\lambda_j) = \frac{P_{\text{signal}}(\lambda_j)}{P_{\text{Noise}}(\lambda_j)} \quad (17)$$

D. Modeling PV of MR Devices in Corona PNoC

We adapt the VARIUS tool [25], similar to [20], to model die-to-die (D2D) as well as within-die (WID) PVs in MRs. We consider photonic devices with a silicon (Si) core and silicon-dioxide (SiO₂) cladding. VARIUS uses a normal distribution to characterize on-chip D2D and WID PVs.

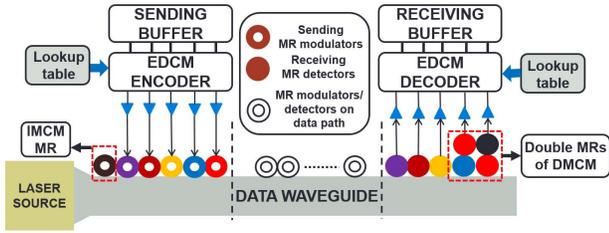


Fig. 3. Overview of cross-layer HYDRA framework that integrates a device-level IM effect-induced signal suppression aware heterodyne crosstalk mitigation (IMCM), a device-level DMCM mechanism, and a circuit-level technique for heterodyne crosstalk mitigation (EDCM).

The key parameters are mean (μ), variance (σ^2), and density (ω) 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 (μ) of wavelength variation of an MR is its nominal resonance wavelength. We consider a DWDM wavelength range in the *C*- and *L*-bands [15], with a starting wavelength of 1550 nm and a channel spacing of 0.8 nm. Hence, those wavelengths are the means for each MR modeled. The variance (σ^2) of wavelength variation is determined based on the laboratory fabrication data [8] and our target die size. We consider a 256-core chip with a die size of 400 mm² at a 22-nm process node. For this die size, we consider a WID standard deviation (σ_{WID}) of 0.61 nm [20] and D2D standard deviation (σ_{D2D}) of 1.01 nm [20]. We also consider a density (ω) of 0.5 [20] for this die size, which is the parameter that determines the range of WID spatial correlation required by the VARIUS tool. With these parameters, we use VARIUS to generate 100 PV maps, these maps are used to model PV in Corona PNoC.

IV. HYDRA FRAMEWORK: OVERVIEW

Our proposed cross-layer HYDRA framework enables crosstalk resilience in DWDM-based PNoC architectures by integrating device-level and circuit-level enhancements that seamlessly work together. Fig. 3 gives a high-level overview of our framework. The IM effect-induced signal suppression aware heterodyne crosstalk mitigation (IMCM) scheme employs additional MRs to decrease wavelength-specific crosstalk noise at the detectors of DWDM-based photonic links. The DMCM mechanism employs DMRs as signal filters to reduce the crosstalk noise at the detectors. This technique improves OSNR in DWDM-based photonic links. However, excessive usage of DMRs (or higher-order filters) increases area, PV redress power (static power required to counter PV-induced resonance drifts in the DMRs), and laser power overheads for PNoC architectures [46]. Thus, to reduce these overheads, we also devise a circuit-level technique for heterodyne crosstalk mitigation (EDCM) that uses a 5-bit encoding mechanism to intelligently reduce undesirable data value occurrences in a photonic waveguide. This allows for further reduction in crosstalk noise and more effectively improves OSNR in DWDM-based PNoC architectures. The following three sections present details of the IMCM, DMCM, and EDCM techniques.

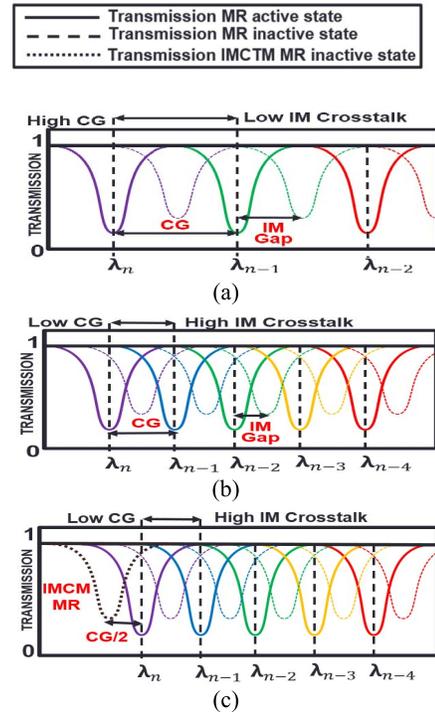


Fig. 4. Transmission spectrum of MR groups with (a) high CG, (b) low CG, and (c) IMCM at low CG.

V. CROSSTALK MITIGATION WITH IMCM

A. IM Crosstalk Analysis

IM crosstalk occurs when the neighboring MR modulators modulate a resonance wavelength of an MR modulator. As discussed earlier in Sections I and III-A, this IM crosstalk suppresses the wavelength signals. The amount of signal suppression depends on the CG between two adjacent wavelengths and IM gap, which is the gap between resonance wavelengths of an MR in active and inactive states. Fig. 4 shows the transmission spectrum of MR groups with high and low CG and IM gap. As the CG and IM gap reduces from low density DWDM [Fig. 4(a)] to high density DWDM [Fig. 4(b)], IM crosstalk increases, as is evident from the intersection of the transmission spectrum of inactive MRs with wavelengths in the waveguide ($\lambda_1 - \lambda_n$). As a result, the suppression of signal power also increases. This IM crosstalk-induced suppression (or loss) in the signal power for given CG and IM gap values is modeled using the factor γ given in (6). Note that for our study, we set the CG and IM gap values to be 0.8 and 0.4 nm, respectively.

All MR modulators at a node in a DWDM waveguide have neighbors on both sides except the first and last modulators. So, compared to all other wavelengths in the DWDM spectrum, the first (λ_1) and last (λ_n) wavelengths are less suppressed (suffer less loss) due to the IM crosstalk caused by the neighboring MRs. Moreover, the waveguide propagation loss incurred on a signal, when it passes by a modulator node, is negligible due to the negligible propagation length (a few hundred micrometers) of the node. Hence, neglecting the waveguide propagation losses incurred by the modulator

node, the first and last wavelengths of the DWDM spectrum have the lowest signal power losses (corresponding to IM crosstalk induced signal power suppression) and the highest signal strengths. When the modulated DWDM wavelengths (signals) travel along a photonic waveguide to the target receiver node (where multiple receivers corresponding to the DWDM wavelengths are arrayed along the waveguide), both λ_1 and λ_n still have high signal strengths compared to all other wavelengths, as all modulated wavelengths propagate the same distance from the modulator node to the target receiver node. Note that each receiver at a receiver node is comprised of a filter MR that is resonant to a specific wavelength and a detector at the drop port of this filter MR.

At a receiver node, the first (λ_1) wavelength signal gets filtered and detected by the first receiver (filtered by the constituent filter MR and detected by the detector). As a result, the signal strength of the first wavelength becomes negligible in the waveguide after the first receiver. Now, from (13), the crosstalk noise power at a detector depends on the power (or strengths) of the neighboring signals at the corresponding detector-coupled MR filter. Therefore, the negligible signal strength of the first wavelength after the first receiver does not add crosstalk noise in the detectors of the succeeding neighboring receivers. In contrast, the last (λ_n) wavelength, which also has higher signal strength, gets filtered by the last receiver in the receiver node. This last (λ_n) wavelength signal has to pass by all receivers along the waveguide before being filtered and detected. On its way to the detector of the last receiver, λ_n incurs crosstalk noise in the detectors of all receivers in the receiver node depending on the relative detuning of the resonance passbands of the receiver's filter MRs from the modulated λ_n signal spectrum. As the strength of this λ_n signal is high, the incurred crosstalk noise is also high.

B. IM-Aware Crosstalk Mitigation

Based on the observations in Section V-A, we propose an IM effect-induced signal suppression aware heterodyne crosstalk mitigation (IMCM) scheme to decrease crosstalk noise in the detectors of DWDM-based photonic links. In IMCM, to reduce signal strength of the last wavelength in the DWDM spectrum before it reaches the target receiver node, we propose placing an additional MR at each modulating and receiving node. This extra MR is tuned near to the last (λ_n) wavelength of DWDM with a tuning distance of half the CG (CG/2) of the DWDM spectrum [as shown in Fig. 4(c)]. This extra MR increases signal suppression of this last (λ_n) wavelength. Thus, it creates uniform signal strength across all wavelengths used in the DWDM spectrum. As a result, the heterodyne noise produced by λ_n in its neighboring receivers is reduced to about the same level as the noise produced by all other wavelengths in their respective neighbor receivers. Thus, as this mechanism reduces crosstalk noise levels without reducing the relative signal strengths, it results in improved OSNR at the detectors. To implement the IMCM technique in PNoCs, there is a need to increase the number of MRs in all modulating and receiving nodes by one. The increase of MRs along the waveguides

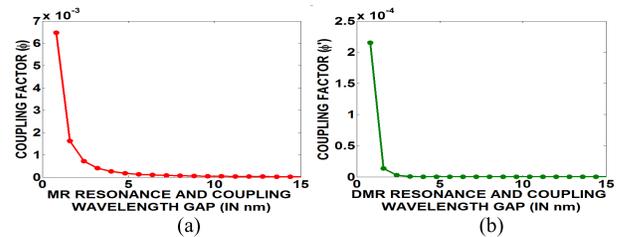


Fig. 5. Coupling factor (ϕ/ϕ') variation with increase in gap between the nonresonant wavelength available in the photonic waveguide, and the resonance wavelength of (a) single MR filter and (b) DMR filter.

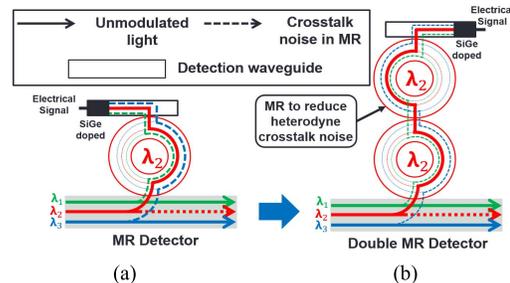


Fig. 6. Crosstalk mitigation with DMRs. (a) MR detector operation when receiving its resonance wavelength. (b) DMR operation when receiving its resonance wavelength.

increases through loss and laser power and this overhead in our analysis in Section IX.

VI. CROSSTALK MITIGATION WITH DMCM

Crosstalk noise in the detectors of DWDM-based PNoCs is mainly caused due to inefficient coupling of filter MRs, as filter MRs in their active mode couple not only photonic power from their resonance wavelengths but also a small amount of photonic power from other wavelengths in the waveguide. The coupling factor ϕ in Fig. 5(a) represents the fraction of signal power of nonresonant wavelength coupled by an MR filter. This coupled power is then dropped on a detector at the MR's drop port. Fig. 5(a) illustrates the variation of ϕ [using (6)] with increase in gap between the MR resonance wavelength and the nonresonant wavelength available in the waveguide. It can be seen that ϕ decreases abruptly with an increase in this gap. The first immediate nonresonance wavelength has almost $4\times$ higher coupling factor than the second immediate nonresonance wavelength, considering a channel spacing of 0.8 nm, $Q = 12\,500$, and 5-Gbps bitrate. We choose these values of channel spacing, Q , and bitrate, as they provide optimal value of total filter penalty for single MR filters (as projected from [41]). Thus nonresonant wavelengths closer to the MR filter's resonance wavelengths create greater cross-talk noise.

One way of reducing this crosstalk noise is to increase the Q-factor of MR filters so that ϕ is reduced. But doing so would increase the photon-lifetime in MR filters limiting their maximum allowable bitrate [42]. An alternate method for reducing crosstalk is to use the second-order filters with DMRs, as used in [46] and [47], for steeper roll-off of filter response. The use of a DMR filter in place of a single MR

filter is depicted in Fig. 6. To further reduce crosstalk, use of filter MRs of even higher order (third order or higher) is possible, but as explained in [46], the use of higher-order MR filters and the choice of Q for the MR stages trade off crosstalk suppression with signal degradation due to signal sidelobe truncation. From [46], the DMRs present lower signal degradation power penalty than the third-order and first-order (single MR) MR filters. The optimal crosstalk performance for DMRs is achieved at 12.5-Gbps bitrate or lower with 0.8-nm channel spacing and the individual MRs having the Q -factor of 8000 [46]. For these reasons, in this paper, we use DMR filters with individual MR Q -factor of 8000 to reduce crosstalk noise.

A. Modeling of DMR Filters

In this section, we model the resultant coupling factor ϕ' and signal suppression/loss factor γ' due to the steeper roll-off of a DMR filter response. From [48], in analogy to electronic filter design, the effect of steeper roll-off of a DMR filter response can be modeled as a maximally flat Butterworth filter response. From [48], the shape, and hence the Q -factor of the Butterworth filter response does not change for higher-order filters (and hence for a DMR) except that the roll-off becomes steeper. Therefore, a Butterworth type of DMR filter response can be modeled by simply setting the exponent of the term $2Q'(\lambda_i - \lambda_j)/\lambda_j$ in (6) and (7) to four instead of two. As a result, (6) and (7) can be revised for a DMR to be, respectively

$$\Phi'(\lambda_i, \lambda_j, Q') = \left(1 + \left(\frac{2Q'(\lambda_i - \lambda_j)}{\lambda_j}\right)^4\right)^{-1} \quad (18)$$

$$\gamma'(\lambda_i, \lambda_j, Q') = \left(1 + \left(\frac{2Q'(\lambda_i - \lambda_j)}{\lambda_j}\right)^4\right)^{-1}. \quad (19)$$

Here, as the Q -factor for a Butterworth DMR filter does not change from the Q -factor of a single MR filter, Q' in (18) and (19) can be modeled as the loaded Q -factor of the individual MRs using (4) and (5). We modeled a DMR with an original Q -factor of 8000 [corresponding to self-coupling coefficients $r_1 = 0.985$, $r_2 = 0.985$, and field-transmission coefficient a of 0.985 in (4) and (5)]. Based on this model, we simulated ϕ' using (18). Fig. 5(b) illustrates the variation of ϕ' [using (18)] with increase in gap between the DMR resonance wavelength and the nonresonant wavelength available in the waveguide. By comparing Fig. 5(a) with Fig. 5(b), it is evident that ϕ' of the MR's immediate nonresonant wavelength (with a CG of 0.8 nm) for the DMR filter is about $30\times$ smaller than ϕ for the single MR filter. Since the coupling factor is used to determine crosstalk noise power in the filter-coupled detectors, it is evident that the DMR filter reduces the crosstalk noise power by about $30\times$. Thus, it can be concluded that the use of DMR filters in place of signal MR filters at the receiver nodes of PNOcs results in significantly less crosstalk noise power at the detectors. Thus, our DMCM scheme uses DMR filters in the place of single MR filters and achieves significant reduction in crosstalk noise power and improvement in OSNR at the detectors.

B. Overhead Analysis for our DMCM Scheme

In this section, we discuss the overhead of using DMR filters. From [27] and [48], as depicted in Fig. 6(b), in a DMR, both the constituent MRs should be in resonance with the same wavelength [i.e., λ_2 in Fig. 6(b)] to achieve a smooth filter response without any ripples or multiple peaks. However, in reality, due to the presence of PV, the constituent MRs end up having different resonance wavelengths after fabrication, which results in multiple peaks in the DMR filter's response. Therefore, the resonances of both the individual MRs of a DMR need to be aligned with trimming or tuning, which almost doubles the required trimming or tuning power for DMR filters compared to single-MR filters. In addition to this, a DMR filter incurs crosstalk-induced signal impairment-related power penalty of 0.5 dB for 0.8-nm channel spacing [46] and incurs about 1.5-dB insertion loss [46]. Moreover, thermal stabilization of a DMR requires 0.9 mW more power [46]. Because of all these penalties, too much use of DMR filters result in a very high-power overhead. Nevertheless, we propose an intelligent method of using few DMR filters along with a data encoding mechanism to limit the use and overheads of DMRs and further mitigate the crosstalk noise power in the detectors. This encoding method is presented next.

VII. CROSSTALK MITIGATION WITH EDCM

The crosstalk noise in a detector is also highly dependent on the strengths of the nonresonant signals at the detector. Crosstalk noise increases with increase in signal power of nonresonant wavelengths. Based on this observation, one can conjecture that crosstalk noise may be mitigated by placing one or more "0"s adjacent to "1"s in the data in the waveguide, to reduce photonic signal strength of nonresonant wavelengths. In this section, we present a novel technique (i. e., EDCM) at the circuit level for mitigation of crosstalk noise in DWDM-based PNOcs.

DMRs in the DMCM technique presented in Section VI increase laser power (because of higher signal loss due to higher crosstalk power penalty and insertion loss) and redress power dissipation overheads. These power overheads increase with an increase in the number of DMRs, hence there is a need to reduce the number of DMRs used with photonic waveguides. In DMCM, DMRs are beneficial when there are consecutive "1"s in the parallel data word being transmitted, because consecutive "1"s imply higher signal strength in the immediate nonresonant wavelengths. One way to reduce the number of DMRs while still minimizing the crosstalk noise due to consecutive "1"s is by reducing the number of consecutive "1"s in the parallel data word being transmitted. To do so, we propose a circuit-level scheme that employs a sophisticated encoding mechanism.

Our proposed circuit-level technique for heterodyne crosstalk mitigation (EDCM) places one or more "0"s adjacent to "1"s in the data to restrict the number of consecutive "1"s in the data stream to three. EDCM employs 5-bit encoding for every 4-bit data block to restrict the number of consecutive "1"s to two in the data block, which in turn limits the worst

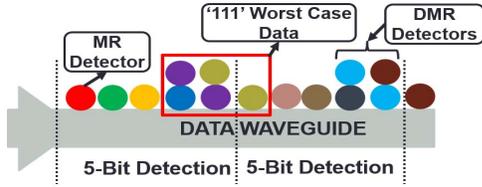


Fig. 7. Organization of MR and DMR detectors in a detecting node on a photonic data waveguide with the EDCM mechanism.

TABLE III
CODE WORDS FOR EDCM TECHNIQUE

Data Word	Code Word	Data Word	Code Word
0000	00000	1000	01000
0001	00001	1001	01001
0010	00010	1010	01010
0011	00011	1011	01011
0100	00100	1100	10100
0101	00101	1101	10010
0110	10011	1110	10001
0111	10101	1111	10000

case number of consecutive “1”s in the data stream to three. Fig. 7 shows the organization of MRs and DMRs in the implementation of the proposed EDCM encoding mechanism along with the location of occurrence of worst case consecutive “1”s. Table III shows the 5-bit codes in the EDCM scheme, to replace 4-bit data words. To implement this encoding technique on a 64-bit word, 16 additional bits are required, which in turn increases the number of MR devices by 25%. However, EDCM reduces the number of DMR detectors required by DMCM and reduces the total number of MR detectors by 12.5%. We propose to use an SRAM-based lookup table with a size of 80-bits to facilitate encoding and decoding of data in each modulating and detecting node for our EDCM mechanism. This encoding and decoding mechanisms incur a delay overhead of approximately one clock cycle, which we account for in our simulation analysis.

VIII. HYDRA INTEGRATION WITH PNOCS

A. Corona PNoC With HYDRA Framework

In this section, we extend the PV-aware crosstalk models of the Corona PNoC from Section III-C to devise PV-aware crosstalk models for Corona enhanced with the HYDRA framework. To integrate HYDRA with the Corona PNoC, we increase the DWDM degree in the MWSR waveguides from 64 to 65 (i.e., channel spacing is reduced from 0.8 to 0.79 nm) and increase the number of MWSR waveguides in each channel from 4 to 5 to facilitate simultaneous transfer of an entire packet (which requires 512 bits before encoding). To distribute optical power between these waveguides, there is also a need to replace 1×4 splitters with 1×5 splitters with a splitting factor of R_{S15} . Because of the increase in DWDM from 64 to 65, the number of modulators in the modulating node increases from 64 to 65. Furthermore, we need to add an additional IMCM MR in all modulating nodes on each MWSR waveguide, thus the

total number of modulators in each modulating node on each MWSR waveguide increases to 66. In the detecting node, first we need to increase the number of detector MRs on each data waveguide from 64 to 65 and second, as shown in Fig. 7, in each group of 5 consecutive detector MRs, we need to replace the last two detector MRs with DMR detectors (replace ϕ , and γ with ϕ' , and γ' , respectively). Therefore, (12)–(15) for worst case signal and crosstalk noise power are changed to the following equations, respectively:

$$P_{\text{signal}}(\lambda_j) = \Phi'(\lambda_j, \lambda_j, Q'_{(63 \times 66) + j}) P_S(\lambda_j, \lambda_j) \quad (20)$$

$$P_{\text{noise}}(\lambda_j) = \sum_{i=1}^n \Phi'(\lambda_i, \lambda_j, Q'_{(63 \times 66) + j}) (P_S(\lambda_i, \lambda_j)) (i \neq j) \quad (21)$$

$$K(\lambda_i) = (R_{S15})(L_{S15})(L_P)^L (L_B)^B \times \prod_{n=1}^{63} \prod_{j=1}^{66} \gamma'(\lambda_i, \lambda_j, Q'_{((n-1) \times 66) + j}) \quad (22)$$

$$\psi(\lambda_i, \lambda_j) = \prod_{k=1}^{(k-1) < j} \gamma'(\lambda_i, \lambda_k, Q'_{(63 \times 66) + k}). \quad (23)$$

B. Firefly PNoC With HYDRA framework

To investigate the efficacy of integrating our HYDRA framework into other PNoC architectures, we integrated it with the Firefly [4] crossbar-based PNoC architecture. Firefly PNoC, for a 256-core system, has eight clusters (C1–C8) with 32 cores in each cluster. Within each cluster, a group of four cores are connected to a router through a concentrator. Thus, each cluster has eight routers (R1–R8) and these routers are electrically connected using a mesh topology. Firefly uses photonic signals for intercluster communication. Unlike the MWSR waveguides used in the Corona crossbar, Firefly uses reservation-assisted single write multiple reader (R-SWMR) data waveguides in its crossbar. Each data channel in Firefly consists of eight SWMR waveguides, with 64 DWDM in each waveguide. Firefly uses only 1/8th of the MRs on each data waveguide compared to Corona, as only eight nodes (one per cluster) are capable of accessing each SWMR waveguide.

In our implementation of Firefly, we considered a power waveguide similar to that used in Corona and determined that the worst case power loss node (N_{WCPL}) is at the detectors of C4R0, which is the router-0 (R0) of cluster-4 (C4) in this architecture. Similar to Corona, in Firefly, the worst case signal and noise power in the detectors of router C4R0 are calculated using (12)–(17) presented in Section III-C. But as Firefly has fewer number of MRs in its data channels, this in turn changes the signal and crosstalk noise power losses before the detector group of N_{WCPL} .

To integrate HYDRA with the Firefly PNoC, we need to increase the DWDM degree in SWMR waveguides from 64 to 65 and increase the number of SWMR waveguides in each channel from 8 to 10 to facilitate simultaneous transfer of an entire packet (which requires 512 bits before encoding). To deal with the increase in DWDM degree, we need to increase the number of modulators and detectors

from 64 to 65 on each SWMR waveguide in a modulating node and detecting node, respectively. Further, we need to add an additional IMCM MR in all modulating and detecting nodes on each SWMR waveguide, which increases the total number of MRs in each modulating and detecting node on each SWMR waveguide to 66. Also, in each detecting node, for each group of five consecutive detector MRs (excluding the IMCM MR in that detecting node), we need to replace the last two detector MRs with DMR detectors (see Fig. 7). Last, we determine worst case OSNR using (20)–(23) with modified through losses.

C. Flexishare PNoC With HYDRA Framework

We also investigated integrating HYDRA with the Flexishare [4], PNoC architecture with 256 cores. We considered a 64-radix, 64-node Flexishare architecture with four cores in each node having 32 data channels for internode communication. Each data channel in Flexishare has four multiple write multiple read (MWMR) waveguides with 64 DWDM in each waveguide. Similar to the MWSR data waveguides of Corona, MWMR data waveguides in Flexishare also uses the models from (12)–(17) presented in Section III-C, to determine the received crosstalk noise and OSNR at detectors for the node with worst case power loss (N_{WCPL}), which corresponds to detectors of node 63 (R_{63}).

To integrate HYDRA with Flexishare, we need to increase the DWDM degree in the MWMR waveguides from 64 to 65 and increase the number of MWMR waveguides in each channel from 4 to 5 to simultaneously transfer 512 bits. We also need to increase the number of modulators and detectors from 64 to 65 on each MWMR waveguide in each modulating and detecting node. Similar to the Firefly PNoC, we need to add an additional IMCM MR in all modulating and detecting nodes on each MWMR waveguide, which increases the total number of MRs in each modulating and detecting node on each SWMR waveguide to 66, respectively. In the detecting nodes of Flexishare, for each group of five consecutive detector MRs (excluding the IMCM MR in that detecting node), we need to replace the last two detector MRs with DMR detectors. Last, we can use (20)–(23) to determine worst case OSNR.

IX. EVALUATION

A. Simulation Setup

To evaluate the efficacy of our proposed cross-layer crosstalk noise mitigation framework HYDRA which combines device layer (IMCM, DMCM) and circuit layer (EDCM) mechanisms for DWDM-based PNoCs, we integrate the framework with the Corona, Firefly, and Flexishare crossbar-based PNoCs, as explained in Section VIII. We modeled and performed simulation-based analysis of the HYDRA-enhanced Corona, Firefly, and Flexishare PNoCs using a cycle-accurate SystemC-based NoC simulator, for a 256-core single-chip architecture at 22 nm. We validated the simulator in terms of power dissipation and energy consumption based on the results obtained from the DSENT tool [30]. We used real-world traffic from applications in the PARSEC benchmark suite [28].

GEM5 full-system simulation [29] of parallelized PARSEC applications was used to generate traces that were fed into our cycle-accurate NoC simulator. We set a “warmup” period of 100 million instructions and then captured traces for the subsequent 1 billion instructions. These traces are extracted from parallel regions of execution of PARSEC benchmark applications. We performed geometric calculations for a 20 mm × 20 mm chip size, to determine lengths of MWSR, SWMR, and MWMR 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 eight cycles at 5-GHz clock frequency [1], [2]. We use a 512-bit packet size, as advocated in the Corona, Firefly, and Flexishare PNoCs.

The static and dynamic energy consumptions of electrical routers and concentrators in Corona, Firefly, and Flexishare PNoCs are based on results from the open source DSENT tool [30]. We model and consider area, power, and performance overheads for our framework implemented with the Corona, Firefly, and Flexishare PNoCs, as follows. HYDRA with Corona, Firefly, and Flexishare PNoCs has an electrical area overhead estimated to be 6.4, 12.7, and 3.4 mm², respectively, and power overheads of 0.23, 0.44, and 0.36 W, respectively, using gate-level analysis and the CACTI 6.5 [31] tool for memory and buffers. The photonic area overheads of Corona, Firefly, and Flexishare architecture are 9.63, 19.83, and 5.2 mm², respectively, based on the physical dimensions [15] of their waveguides, MRs, and splitters. For energy consumption of photonic devices, we adapt model parameters from recent work [6], [32], and [33], with 0.42 pJ/bit for every modulation and detection event and 0.18 pJ/bit for the driver circuits of modulators and photodetectors. The MR trimming power is set to 130 μW/nm [11] for current injection (blue shift) and tuning power is set to 240 μW/nm [12] for heating (red shift).

B. Worst Case OSNR Comparison for Various PNoCs

Our first set of simulation results compares the baseline (without any crosstalk-mitigating enhancements) Corona, Firefly, and Flexishare PNoCs with four variants of these architectures corresponding to three crosstalk-mitigation strategies from prior work (PCTM5B and PCTM6B from [13], PICO from [35]) and our proposed HYDRA framework from this paper. PCTM5B and PCTM6B are encoding schemes that replace each 4-bits of a data word with 5- and 6-bit code words, respectively. These schemes aim to reduce photonic signal strength of immediate nonresonant wavelengths (adjacent wavelengths in DWDM) to decrease crosstalk and improve OSNR in MR detectors. PICO is a process-variation aware crosstalk-mitigation mechanism which also encodes data to reduce photonic signal strength of immediate nonresonant wavelengths based on the PV profile of the receiving MR detectors.

Utilizing the models presented in Sections III and VIII, we calculate the received crosstalk noise and OSNR at detectors for the node with worst case power loss (N_{WCPL}), which correspond to MR detectors in cluster 64 for the Corona

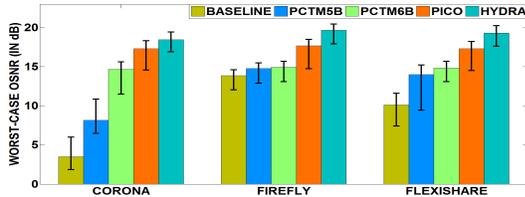


Fig. 8. Worst case OSNR comparison of HYDRA with PCTM5B [5], PCTM6B [5], and PICO [35] for Corona, Firefly, and Flexishare PNoCs. Bars show mean values of worst case OSNR across 100 PV maps; confidence intervals show variation in worst case OSNR.

PNoC, MR detectors of router C4R0 for the Firefly PNoC, and MR detectors of node R_{63} for the Flexishare PNoC. Fig. 8 summarizes the worst case OSNR results for the baseline, PCTM5B, PCTM6B, PICO, and HYDRA configurations of the three PNoC architectures considered. From Fig. 8, it can be observed that Corona PNoC with HYDRA has 5.3 \times , 2.26 \times , 1.25 \times , and 1.06 \times , Firefly PNoC with HYDRA has 1.42 \times , 1.33 \times , 1.32 \times , and 1.13 \times , and Flexishare PNoC with HYDRA has 1.96 \times , 1.41 \times , 1.33 \times , and 1.14 \times worst case OSNR improvements on average, compared to the baseline and PCTM5B, PCTM6B, and PICO enhanced variants of these architectures, respectively.

Both PCTM5B and PCTM6B eliminate occurrences of “111” in a data word and have limited occurrences of “11,” which helps to reduce crosstalk noise in the detectors. But these techniques are unable to eliminate all occurrences of “11,” because of which these techniques are unable to achieve higher reduction in crosstalk noise and significant improvement in OSNR. PICO considers the PV-profile of detecting nodes and performs encoding on specific wavelengths where there is a high signal loss due to trimming to reduce crosstalk noise and improves OSNR in PNoCs. But even with the PICO technique, there still exist occurrences of “111” and “11,” because of which OSNR gains with PICO are on the lower side. In contrast, HYDRA virtually eliminates all of the occurrences of “111” and “11” from the data word by combining benefits from IMCM and DMCM, and using *EDCM*’s 5-bit encoding mechanism. Although *EDCM*’s 5-bit encoding still results in limited occurrences of “111” and “11” in a data word, the DMRs of DMCM reduce the impact of consecutive “1”s in the data word by removing crosstalk noise generated by these “1”s in detector MRs. Thus HYDRA demonstrates higher OSNR gains compared to the best known previously proposed techniques. Furthermore, the OSNR values achieved with HYDRA (see Fig. 8) are sufficient to enable reliable data transfers in PNoCs such as Corona, Firefly, and Flexishare.

C. Overhead Analysis of HYDRA With Various PNoCs

Our last set of results quantifies the overhead for the proposed HYDRA framework and other techniques when used with the Corona, Firefly, and Flexishare PNoCs. Fig. 9(a) and (b) presents detailed simulation results that quantify the average network packet latency and energy-delay product (EDP) for five Corona configurations. Results are shown for 12 multithreaded PARSEC benchmarks. From Fig. 9(a), it can be seen that on average, Corona with HYDRA has 9.24% higher latency compared to the baseline. The additional

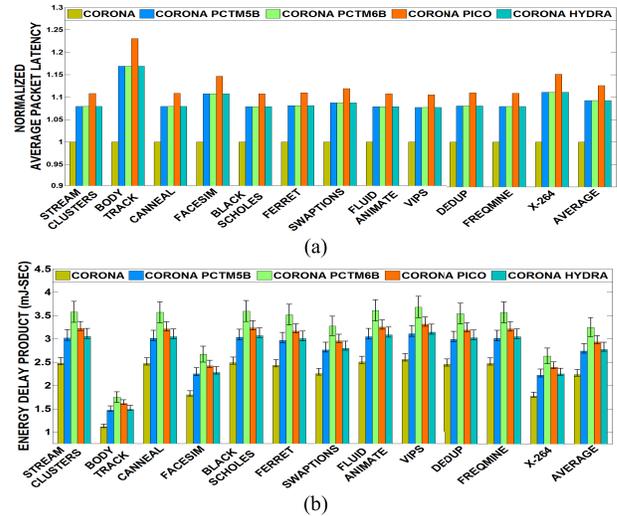


Fig. 9. (a) Normalized average latency. (b) EDP comparison between Corona baseline and Corona configurations with PCTM5B, PCTM6B, PICO, and HYDRA techniques, for PARSEC benchmarks. Latency results are normalized to the baseline Corona results. In the EDP plot, bars represent mean values of EDP across 100 PV maps; confidence intervals show variation in EDP.

delay due to encoding and decoding of data with HYDRA, PCTM5B, PCTM6B, and PICO contributes to their increase in average latency. The penalty due to encoding/decoding is approximately 1 cycle in PCTM5B, PCTM6B, and HYDRA. Thus HYDRA has a similar overhead compared to PCTM5B and PCTM6B. However, *PICO* has a 2-cycle penalty, which increases its delay compared to HYDRA by 3.1%. Note that for the chosen clock frequency, PV in photonic components does not change the number of clock cycles for various operations, such as encoding/decoding, modulation/detection etc., therefore, Fig. 9(a) does not have confidence intervals or variations in packet latency due to PV.

From the results for EDP shown in Fig. 9(b), it can be seen that on average, the Corona configuration with our HYDRA framework has 24.3% higher EDP compared to the baseline. The increase in EDP for Corona with HYDRA is not only due to the increase in average latency, but also due to the addition of extra bits for encoding and decoding, which leads to an increase in the amount of photonic hardware in the architectures (more number of MRs, complex splitters). This in turn increases static power dissipation. Dynamic power also increases in these architectures, but by much less amount. However, EDP for Corona with HYDRA is 17.1% and 5.7% lower compared to PCTM6B and PICO, respectively. The higher latency of PICO compared to HYDRA increases its EDP, whereas HYDRA has lower EDP compared to PCTM6B because HYDRA conserves laser and MR trimming/tuning power due to a lower photonic hardware footprint compared to PCTM6B. The EDP for Corona with HYDRA is 1.3% higher compared to PCTM5B. Although PCTM5B and HYDRA have similar average latency, the increase in number of MRs in HYDRA due to the presence of IMCM MRs and DMRs increases its laser and trimming/tuning power, which in turn increases its EDP.

Fig. 10(a) and (b) summarizes the average network packet latency and EDP results for the five configurations of

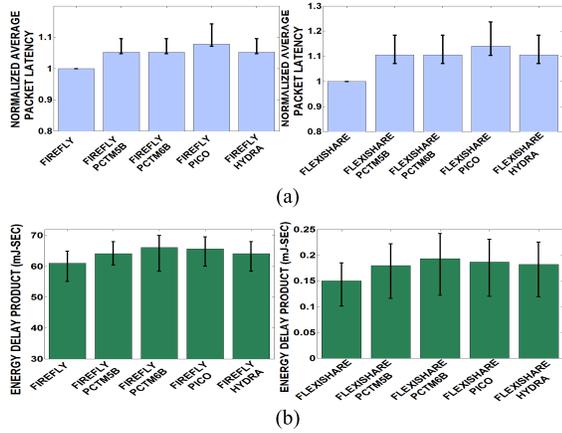


Fig. 10. (a) Normalized average latency. (b) EDP comparison between different variants of Firefly and Flexishare PNoCs, which include their baselines and their variants with PCTM5B, PCTM6B, PICO, and HYDRA techniques, for PARSEC benchmark applications. Latency results are normalized with their respective baseline architecture results. Bars represent mean values of average latency and EDP for 100 PV maps; confidence intervals show variation in average latency and EDP across PARSEC benchmarks.

Firefly and Flexishare PNoCs. Results are shown for twelve multithreaded PARSEC benchmarks and are averaged across these benchmark applications, for brevity.

From Fig. 10(a), it can be observed that on average, Firefly with HYDRA has 5.2% and Flexishare with HYDRA has 10.6% higher latency compared to their respective baselines. The additional delay due to encoding and decoding of data with HYDRA contributes to its increase in average latency over the respective baselines of Firefly and Flexishare PNoCs. The latency overhead for Firefly with HYDRA is lower compared to Corona and Flexishare with HYDRA. This is because Firefly is a hybrid PNoC where some portion of data traverses through electrical links. This data over electrical links is unaffected by the extra encoding/decoding delays in HYDRA, whereas in Corona and Flexishare the entire traffic traverses through photonic waveguides. Much like Corona [Fig. 9(a)], the Firefly and Flexishare architectures with HYDRA have similar latency values compared to these architectures with PCTM5B and PCTM6B [Fig. 10(a)]. Furthermore, Firefly with HYDRA has 2.7% and Flexishare with HYDRA has 3.2% lower latency compared to PICO. Reduction in number of encoding or decoding cycles from 2 to 1 from PICO to HYDRA reduces average latency of HYDRA compared to PICO.

From the results for EDP shown in Fig. 10(b), it can be seen that on average, the Firefly and Flexishare configurations with our HYDRA framework have 5% and 22% higher EDP compared to their respective baselines. EDP overhead for Firefly is relatively lower compared to the Corona and Flexishare architectures because of its lower latency overheads and smaller increase in laser/trimming power due to lesser increase in the amount of photonic hardware. Firefly with HYDRA has 3.1% and 2.4%, and Flexishare with HYDRA has 5.9% and 2.4% lower EDP compared to the respective architecture configurations with PCTM6B and PICO. Additionally, compared to Firefly and Flexishare config-

urations with PCTM5B, the configurations of the same architectures with HYDRA framework have 0.6% and 1.5% higher EDP, respectively.

X. CONCLUSION

We have presented a novel cross-layer crosstalk-mitigation framework for the reduction of crosstalk noise in the detectors of DWDM-based PNoC architectures. Our proposed HYDRA framework seamlessly integrates two device layer and a circuit layer technique to enable interesting tradeoffs between reliability, performance, and energy overheads for the Corona, Firefly, and Flexishare crossbar-based PNoC architectures. Our simulation-based analysis shows that the HYDRA framework improves worst case OSNR by up to $5.3\times$ compared to the baseline architectures, and by up to $1.14\times$ compared to the best known PNoC crosstalk-mitigation scheme from prior work. Thus, HYDRA is an attractive solution to enhance reliability in emerging DWDM-based PNoCs.

REFERENCES

- [1] D. Vantrease, N. Binkert, R. Schreiber, and M. H. Lipasti, "Light speed arbitration and flow control for nanophotonic interconnects," in *Proc. IEEE/ACM MICRO*, Dec. 2009, pp. 304–315.
- [2] S. V. R. Chittamuru, S. Pasricha, and S. Desai, "SWIFTNoC: A reconfigurable silicon-photonic network with multicast-enabled channel sharing for multicore architectures," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 13, no. 4, 2017, Art. no. 58.
- [3] Y. Pan, P. Kumar, J. Kim, Y. Zhang, A. Choudhary, and G. Memik, "Firefly: Illuminating future network-on-chip with nanophotonics," in *Proc. ISCA*, 2009, pp. 429–440.
- [4] Y. Pan, J. Kim, and G. Memik, "FlexiShare: Channel sharing for an energy-efficient nanophotonic crossbar," in *Proc. HPCA*, 2010, pp. 1–12.
- [5] S. V. R. Chittamuru, S. Desai, and S. Pasricha, "Reconfigurable silicon-photonic network with improved channel sharing for multicore architectures," in *Proc. ACM GLSVLSI*, May 2015, pp. 63–68.
- [6] L. H. K. Duong *et al.*, "Coherent crosstalk noise analyses in ring-based optical interconnects," in *Proc. DATE*, 2015, pp. 501–506.
- [7] K. Padmaraju, X. Zhu, L. Chen, M. Lipson, and K. Bergman, "Intermodulation crosstalk characteristics of WDM silicon microring modulators," *IEEE Photon. Technol. Lett.*, vol. 26, no. 14, pp. 1478–1481, Jul. 15, 2014.
- [8] S. K. Selvaraja, "Wafer-scale fabrication technology for silicon photonic integrated circuits," Ph.D. dissertation, Dept. Inf. Technol., Ghent Univ., Ghent, Belgium, 2011.
- [9] S. V. R. Chittamuru, I. G. Thakkar, and S. Pasricha, "Process variation aware crosstalk mitigation for DWDM based photonic NoC architectures," in *Proc. ISQED*, Mar. 2016, pp. 57–62.
- [10] Z. Li *et al.*, "Reliability modeling and management of nanophotonic on-chip networks," *IEEE Trans. Very Large Scale Integr. (VLSI) Syst.*, vol. 20, no. 1, pp. 98–111, Jan. 2010.
- [11] J. Ahn *et al.*, "Devices and architectures for photonic chip-scale integration," *Appl. Phys. A*, vol. 95, no. 4, pp. 989–997, 2009.
- [12] C. Sun *et al.*, "Single-chip microprocessor that communicates directly using light," *Nature*, vol. 528, no. 7583, pp. 24–31, 2015.
- [13] S. V. R. Chittamuru and S. Pasricha, "Crosstalk mitigation for high-radius and low-diameter photonic NoC architectures," *IEEE Des. Test*, vol. 32, no. 3, pp. 29–39, Jun. 2015.
- [14] Q. Xu, B. Schmidt, J. Shakya, and M. Lipson, "Cascaded silicon microring modulators for WDM optical interconnection," *Opt. Exp.*, vol. 14, no. 20, pp. 9431–9436, 2006.
- [15] S. Xiao, M. H. Khan, H. Shen, and M. Qi, "Modeling and measurement of losses in silicon-on-insulator resonators and bends," *Opt. Exp.*, vol. 15, no. 17, pp. 10553–10561, 2007.
- [16] L. H. K. Duong *et al.*, "A case study of signal-to-noise ratio in ring-based optical networks-on-chip," *IEEE Design Test*, vol. 31, no. 5, pp. 55–65, Oct. 2014.
- [17] S. V. R. Chittamuru and S. Pasricha, "Improving crosstalk resilience with wavelength spacing in photonic crossbar-based network-on-chip architectures," in *Proc. IEEE MWSCAS*, Aug. 2015, pp. 1–4.

- [18] S. V. R. Chittamuru and S. Pasricha, "SPECTRA: A framework for thermal reliability management in silicon-photonics networks-on-chip," in *Proc. IEEE VLSI*, Jan. 2016, pp. 86–91.
- [19] D. A. Neamen, *Semiconductor Physics and Devices: Basic Principles*, 3rd ed. New York, NY, USA: McGraw-Hill, 2002, pp. 115–139.
- [20] Y. Xu, J. Yang, and R. Melhem, "Tolerating process variations in nanophotonic on-chip networks," in *Proc. ISCA*, Jun. 2012, pp. 142–152.
- [21] D. Dang, S. V. R. Chittamuru, R. Mahapatra, and S. Pasricha, "Islands of heaters: A novel thermal management framework for photonic NoCs," in *Proc. IEEE/ACM ASPDAC*, Jan. 2017, pp. 306–311.
- [22] I. G. Thakkar, S. V. R. Chittamuru, and S. Pasricha, "Run-time laser power management in photonic NoCs with on-chip semiconductor optical amplifiers," in *Proc. IEEE/ACM NOCS*, Aug. 2016, pp. 1–4.
- [23] C. Li *et al.*, "Silicon photonic transceiver circuits with microring resonator bias-based wavelength stabilization in 65 nm CMOS," *IEEE J. Solid-State Circuits*, vol. 49, no. 6, pp. 1419–1436, Jun. 2014.
- [24] R. G. Beausoleil, "Large-scale integrated photonics for high-performance interconnects," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 7, no. 2, 2011, Art. no. 6.
- [25] S. R. Sarangi, B. Greskamp, R. Teodorescu, J. Nakano, A. Tiwari, and J. Torrellas, "VARIUS: A model of process variation and resulting timing errors for microarchitects," *IEEE Trans. Semicond. Manuf.*, vol. 21, no. 1, pp. 3–13, Feb. 2008.
- [26] J. E. Heebner, "Nonlinear optical whispering gallery microresonators for photonics," Ph.D. dissertation, Inst. Opt., Univ. Rochester, Rochester, NY, USA, 2003.
- [27] W. Bogaerts *et al.*, "Silicon microring resonators," *Laser Photon. Rev.*, vol. 6, no. 1, pp. 47–73, 2012.
- [28] C. Bienia, S. Kumar, K. Li, and J. P. Singh, "The PARSEC benchmark suite: Characterization and architectural implications," in *Proc. PACT*, Oct. 2008, pp. 72–81.
- [29] N. Binkert *et al.*, "The gem5 simulator," *ACM SIGARCH Comput. Arch. News*, vol. 39, no. 2, pp. 1–7, May 2011.
- [30] C. Sun *et al.*, "DSENT—A tool connecting emerging photonics with electronics for Opto-electronic networks-on-chip modeling," in *Proc. IEEE/ACM NOCS*, May 2012, 201–210.
- [31] *CACTI 6.5*. Accessed: Dec. 1, 2016. [Online]. Available: <http://www.hpl.hp.com/research/cacti/>
- [32] X. Zheng *et al.*, "Ultra-efficient 10 Gb/s hybrid integrated silicon photonic transmitter and receiver," *Opt. Exp.*, vol. 19, no. 6, pp. 5172–5186, 2011.
- [33] P. Grani and S. Bartolini, "Design options for optical ring interconnect in future client devices," *ACM J. Emerg. Technol. Comput. Syst.*, vol. 10, no. 4, 2014, Art. no. 30.
- [34] K. Padmaraju and K. Bergman, "Resolving the thermal challenges for silicon microring resonator devices," *Nanophotonics*, vol. 3, nos. 4–5, pp. 269–281, 2013.
- [35] S. V. R. Chittamuru, I. G. Thakkar, and S. Pasricha, "PICO: Mitigating heterodyne crosstalk due to process variations and intermodulation effects in photonic NoCs," in *Proc. IEEE/ACM DAC*, Jun. 2016, Art. no. 39.
- [36] P. K. Tien, "Light waves in thin films and integrated optics," *Appl. Opt.*, vol. 10, no. 11, pp. 2395–2413, Nov. 1971.
- [37] G. T. Reed and A. P. Knights, *Silicon Photonics: An Introduction*. Hoboken, NJ, USA: Wiley, 2004, pp. 70–75.
- [38] S. Bahirat and S. Pasricha, "Exploring hybrid photonic networks-on-chip foremerging chip multiprocessors," in *Proc. CODES+ISSS*, 2009, pp. 129–136.
- [39] R. Wu, C.-H. Chen, J.-M. Fedeli, M. Fournier, R. G. Beausoleil, and K.-T. Cheng, "Compact modeling and system implications of microring modulators in nanophotonic interconnects," in *Proc. SLIP*, Jun. 2015, pp. 1–6.
- [40] Lumerical Solutions Inc, Vancouver, BC, Canada. *MODE Solutions*. Accessed: Nov. 1, 2016. [Online]. Available: <http://www.lumerical.com/tead-products/mode/>
- [41] M. Bahadori, D. Nikolova, S. Rumley, C. P. Chen, and K. Bergman, "Optimization of microring-based filters for dense WDM silicon photonic interconnects," in *Proc. IEEE OI*, Apr. 2015, pp. 84–85.
- [42] I. G. Thakkar, S. V. R. Chittamuru, and S. Pasricha, "A comparative analysis of front-end and back-end compatible silicon photonic on-chip interconnects," in *Proc. IEEE/ACM SLIP*, Jun. 2016, pp. 1–8.
- [43] R. Hendry, D. Nikolova, S. Rumley, N. Ophir, and K. Bergman, "Physical layer analysis and modeling of silicon photonic WDM bus architectures," in *Proc. HiPEAC Workshop*, 2014, pp. 20–22.
- [44] B. G. Lee, B. A. Small, K. Bergman, Q. Xu, and M. Lipson, "Transmission of high-data-rate optical signals through a micrometer-scale silicon ring resonator," *Opt. Lett.*, vol. 31, no. 81, pp. 2701–2703, 2006.
- [45] S. V. R. Chittamuru, I. G. Thakkar, and S. Pasricha, "Analyzing voltage bias and temperature induced aging effects in photonic interconnects for manycore computing," in *Proc. IEEE/ACM SLIP*, Jun. 2017, pp. 1–8.
- [46] N. Ophir, C. Mineo, D. Mountain, and K. Bergman, "Silicon photonic microring links for high-bandwidth-density, low-power chip I/O," *IEEE Micro*, vol. 33, no. 1, pp. 54–67, Jan. 2013.
- [47] I. G. Thakkar, S. V. R. Chittamuru, and S. Pasricha, "Mitigation of homodyne crosstalk noise in silicon photonic NoC architectures with tunable decoupling," in *Proc. ACM CODES+ISSS*, Oct. 2016, pp. 1–10.
- [48] D. G. Rabus, *Integrated Ring Resonators: The Compendium*. Berlin, Germany: Springer-Verlag, 2007, pp. 25–34.
- [49] S. Pasricha and S. Bahirat, "OPAL: A multi-layer hybrid photonic NoC for 3D ICs," in *Proc. ASPDAC*, Jan. 2011, pp. 345–350.
- [50] S. Bahirat and S. Pasricha, "METEOR: Hybrid photonic ring-mesh network-on-chip for multicore architectures," *ACM Trans. Embedded Comput. Syst.*, vol. 13, no. 3s, Mar. 2014, Art. no. 116.



Sai Vineel Reddy Chittamuru (S'14) received the B.Tech. degree in electrical engineering from IIT Kharagpur, Kharagpur, India, in 2011. He is currently pursuing the Ph.D. degree with the Electrical and Computer Engineering Department, Colorado State University, Fort Collins, CO, USA.

His current research interests include embedded system, systems on chip, and optical networks-on-chip.

Mr. Chittamuru was a recipient of the Best Paper Awards at SLIP'2016 and GLSVLSI'15.



Ishan G. Thakkar (S'14) received the M.S. degree in electrical engineering from Colorado State University, Fort Collins, CO, USA, where he is currently pursuing the Ph.D. degree in electrical engineering.

His current research interests include DRAM architectures, nonvolatile memories, high-speed optical interfaces, and optical networks on chip.

Mr. Thakkar was a recipient of the Best Paper Award at SLIP'2016.



Sudeep Pasricha (M'02–SM'13) received the Ph.D. degree in computer science from the University of California, Irvine, CA, USA, in 2008.

He is currently a Monfort and Rockwell-Anderson Professor of electrical and computer engineering with Colorado State University, Fort Collins, CO, USA. His current research interests include energy efficiency and fault tolerance for multicore computing.

Dr. Pasricha was a recipient of the Best Paper Awards at SLIP'16, GLSVLSI'15, AICCSA'11, ISQED'10, and ASPDAC'06. He is on the Editorial Board of various journals such as TCAD, TMSCS, and TECS. He currently is or has been with the Organizing Committee and TPC of various conferences such as DAC, DATE, NOCS, ESWEEK, and VLSID.