SOTERIA: Exploiting Process Variations to Enhance Hardware Security with Photonic NoC Architectures

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

ABSTRACT
Photonic networks-on-chip (PNoCs) enable high bandwidth on-chip data transfers by using photonic waveguides capable of dense-wavelength-division-multiplexing (DWDM) for signal traversal and microring resonators (MRs) for signal modulation. A Hardware Trojan in a PNoC can manipulate the electrical driving circuit of its MRs to cause the MRs to snoop data from the neighboring wavelength channels in a shared photonic waveguide. This introduces a serious security threat. This paper presents a novel framework called *SOTERIA*† that utilizes process variation based authentication signatures along with architecture-level enhancements to protect data in PNoC architectures from snooping attacks. Evaluation results indicate that our approach can significantly enhance the hardware security in DWDM-based PNoCs with minimal overheads of up to 10.6% in average latency and of up to 13.3% in energy-delay-product (EDP).

Categories and Subject Descriptors: [Networks] Network on chip; [Security and privacy] Security in hardware; [Hardware] Emerging optical and photonic technologies
General Terms – Security, Performance, Experimentation
Keywords – Process Variations, Hardware Security, Photonic NoCs

1. INTRODUCTION

To cope with the growing performance demands of modern Big Data and cloud computing applications, the complexity of hardware in modern chip-multiprocessors (CMPs) has increased. To reduce the hardware design time of these complex CMPs, third-party hardware IPs are frequently used. But these third party IPs can introduce security risks [1]-[2]. For instance, the presence of Hardware Trojans (HTs) in the third-party IPs can lead to leakage of critical and sensitive information from modern CMPs [3]. Thus, security researchers that have traditionally focused on software-level security are now increasingly interested in overcoming hardware-level security risks.

Many CMPs today use electrical networks-on-chip (ENoCs) for inter-core communication. ENoCs use packet-switched network fabrics and routers to transfer data between on-chip components [4]. Recent developments in silicon photonics have enabled the integration of photonic components and interconnects with CMOS circuits on a chip. Photonic NoCs (PNoCs) provide several prolific advantages over their metallic counterparts (i.e., ENoCs), including the ability to communicate at near light speed, larger bandwidth density, and lower dynamic power dissipation [5]. These advantages motivate the use of PNoCs for inter-core communication in modern CMPs [6].

Several PNoC architectures have been proposed to date (e.g., [7]-[9]). These architectures employ on-chip photonic links, each of which connects two or more gateway interfaces. A gateway interface (GI) connects the PNoC to a cluster of processing cores. Each photonic link comprises one or more photonic waveguides and each waveguide can support a large number of dense-wavelength-division-multiplexed (DWDM) wavelengths. A wavelength serves as a data signal carrier. Typically, multiple data signals are generated at a source GI in the electrical domain (as sequences of logical 1 and 0 voltage levels) which are modulated onto the multiple DWDM carrier wavelengths simultaneously, using a bank of modulator MRs at the source GI [10]. The data-modulated carrier wavelengths traverse a link to a destination GI, where an array of detector MRs filter them and drop them on photodetectors to regenerate electrical data signals.

In general, each GI in a PNoC is able to send and receive data in the optical domain on all of the utilized carrier wavelengths. Therefore, each GI has a bank of modulator MRs (i.e., modulator bank) and a bank of detector MRs (i.e., detector bank). Each MR in a bank resonates with and operates on a specific carrier wavelength. Thus, the excellent wavelength selectivity of MRs and DWDM capability of waveguides enable high bandwidth parallel data transfers in PNoCs.

Similar to CMPs with ENoCs, the CMPs with PNoCs are expected to use several third party IPs/, and therefore, are vulnerable to security risks [11]. For instance, if the entire PNoC used within a CMP is a third-party IP, then this PNoC with HTs within the control units of its GIs can snoop on packets in the network. These packets can be transferred to a malicious core (a core running a malicious program) in the CMP to determine sensitive information.

Unfortunately, MRs are especially susceptible to security threatening manipulations from HTs. In particular, *the MR tuning circuits that are essential for supporting data broadcasts and to counteract MR resonance shifts due to process variations (PV) make it easy for HTs to retune MRs and initiate snooping attacks.* To enable data broadcast in PNoCs, the tuning circuits of detector MRs partially detune them from their resonance wavelengths [8], [12]-[13], such that a significant portion of the photonic signal energy in the data-carrying wavelengths continues to propagate in the waveguide to be absorbed in the subsequent detector MRs. On the other hand, process variations (PV) cause resonance wavelength shifts in MRs [14]. Techniques to counteract PV-induced resonance shifts in MRs involve retuning the resonance wavelengths by using carrier injection/depletion or thermal tuning [6], implemented through MR tuning circuits. An HT in the GI can manipulate these tuning circuits of detector MRs to partially tune the detector MR to a passing wavelength in the waveguide, which enables snooping of the data that is modulated on the passing wavelength. *Such covert data snooping is a serious security risk in PNoCs.*

In this work, we present a framework that protects data from snooping attacks and improves hardware security in PNoCs. Our framework has low overhead and is easily implementable in any existing DWDM-based PNoC without major changes to the architecture. To the best of our knowledge, this is the first work that attempts to improve hardware security for PNoCs. Our novel contributions are:

- We analyze security risks in photonic devices and extend this analysis to link-level, to determine the impact of these risks on PNoCs;
- We propose a circuit-level PV-based security enhancement scheme

DAC '18, June 24–29, 2018, San Francisco, CA, USA

ACM ISBN 978-1-4503-5700-5/18/06…$15.00
https://doi.org/10.1145/3195970.3196118

†*Soteria is the Greek goddess of safety and deliverance from harm*

that uses PV-based authentication signatures to protect data from snooping attacks in photonic waveguides;

- We propose an architecture-level reservation-assisted security enhancement scheme to improve security in DWDM-based PNoCs;
- We combine the circuit- and architecture-level schemes into a holistic framework called *SOTERIA*; and analyze it on the Firefly [8] and Flexishare [9] crossbar-based PNoC architectures.

2. RELATED WORK

Several prior works [11], [16], [17] discuss the presence of security threats in ENoCs and have proposed solutions to mitigate them. In [11], a three-layer security system approach was presented by using data scrambling, packet certification, and node obfuscation to enable protection against data snooping attacks. A symmetric-key based cryptography design was presented in [16] for securing the NoC. In [17], a framework was presented to use permanent keys and temporary session keys for NoC transfers between secure and non-secure cores. *However, no prior work has analyzed security risks in photonic devices and links; or considered the impact of these risks on PNoCs.*

Fabrication-induced PV 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 thermal tuning or localized trimming [6]. Trimming can induce blue shifts in the resonance wavelengths of MRs using carrier injection into MRs, whereas thermal tuning can induce red shifts in MR resonances through heating of MRs using integrated heaters. *To remedy PV, the use of device-level trimming/tuning techniques is inevitable; but their use also enables partial detuning of MRs that can be used to snoop data from a shared photonic waveguide.* In addition, prior works [18]-[19] discuss the impact of PV-remedial techniques on crosstalk noise and proposed techniques to mitigate it. *None of the prior works analyze the impact of PV-remedial techniques on hardware security in PNoCs.*

Our proposed framework in this paper is novel as it enables security against snooping attacks in PNoCs for the first time. Our framework is network agnostic, mitigates PV, and has minimal overhead, while improving security for any DWDM-based PNoC architecture.

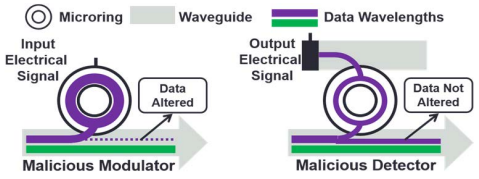

Fig. 1: Impact of (a) malicious modulator MR, (b) malicious detector MR on data in DWDM-based photonic waveguides.

3. HARDWARE SECURITY CONCERNS IN PNOCS

3.1 Device-Level Security Concerns

Process variation (PV) induced undesirable changes in MR widths and heights cause "shifts" in MR resonance wavelengths, which can be remedied using localized trimming and thermal tuning methods. The localized trimming method injects (or depletes) free carriers into (or from) the Si core of an MR using an electrical tuning circuit, which reduces (or increases) the MR's refractive index owing to the electro-optic effect, thereby remedying the PV-induced red (or blue) shift in the MR's resonance wavelength. In contrast, thermal tuning employs an integrated micro-heater to adjust the temperature and refractive index of an MR (owing to the thermo-optic effect) for PV remedy. Typically, the modulator MRs and detectors use the same electro-optic effect (i.e., carrier injection/depletion) implemented through the same electrical tuning circuit as used for localized trimming, to move in and out of resonance (i.e., switch ON/OFF) with a wavelength [19]. *A HT can manipulate this electrical tuning circuit, which may lead to malicious operation of modulator and detector MRs, as discussed next.*

Fig. 1(a) shows the malicious operation of a modulator MR. A malicious modulator MR is partially tuned to a data-carrying wavelength (shown in purple) that is passing by in the waveguide. The malicious modulator MR draws some power from the data-carrying wavelength, which can ultimately lead to data corruption as optical '1's in the data can lose significant power to be altered into '0's. Alternatively, a malicious detector (Fig. 1(b)) can be *partially* tuned to a passing data-carrying wavelength, to filter only a small amount of its power and drop it on a photodetector for data duplication. This small amount of filtered power does not alter the data in the waveguide so that it continues to travel to its target detector for legitimate communication [12]. Thus, malicious detector MRs can snoop data from the waveguide without altering it, which is a major security threat in photonic links. Note that malicious modulator MRs only corrupt data (which can be detected) and do not covertly duplicate it, and are thus not a major security risk.

3.2 Link-Level Security Concerns

Typically, a photonic link is comprised of one or more DWDM-based photonic waveguides. A DWDM-based photonic waveguide uses a modulator bank (a series of modulator MRs) at the source GI and a detector bank (a series of detector MRs) at the destination GI. DWDM-based waveguides can be broadly classified into four types: single-writer-single-reader (SWSR), single-writer-multiple-reader (SWMR), multiple-writer-single-reader (MWSR), and multiple-writer-multiple-reader (MWMR). As SWSR, SWMR, and MWSR waveguides are subsets of an MWMR waveguide, and due to limited space, we restrict our link-level analysis to MWMR waveguides only.

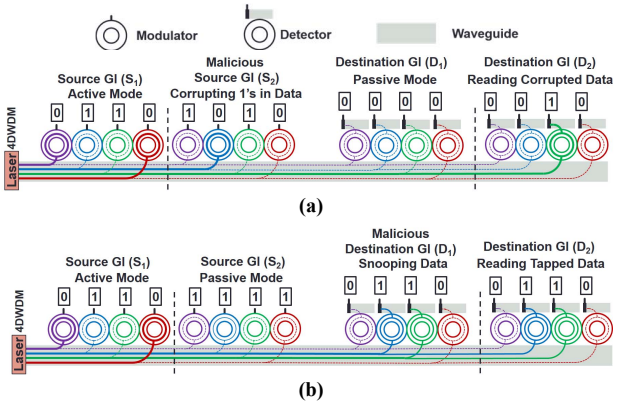

Fig. 2: Impact of (a) malicious modulator (source) bank, (b) malicious detector bank on data in DWDM-based photonic waveguides.

An MWMR waveguide typically passes through multiple GIs, connecting the modulator banks of some GIs to the detector banks of the remaining GIs. Thus, in an MWMR waveguide, multiple GIs (referred to as source GIs) can send data using their modulator banks and multiple GIs (referred to as destination GIs) can receive (read) data using their detector banks. Fig. 2 presents an example MWMR waveguide with two source GIs and two destination GIs. Fig. 2(a) and 2(b), respectively, present the impact of malicious source and destination GIs on this MWMR waveguide. In Fig. 2(a), the modulator bank of source GI S_1 is sending data to the detector bank of destination GI D_2. When source GI S_2, which is in the communication path, becomes malicious with an HT in its control logic, it can manipulate its modular bank to modify the existing '1's in the data to '0's. This ultimately leads to data corruption. For example, in Fig. 2(a), S_1 is supposed to send '0110' to D_2, but because of data corruption by malicious GI S_2, '0010' is received by D_2. Nevertheless, this type of data corruption can be detected or even corrected using parity or error correction code (ECC) bits in the data. Thus, malicious source GIs do not cause major security risks in DWDM-based MWMR waveguides.

Let us consider another scenario for the same data communication path (i.e., from S_1 to D_2). When destination GI D_1, which is in the communication path, becomes malicious with an HT in its control

logic, the detector bank of D_1 can be partially tuned to the utilized wavelength channels to snoop data. In the example shown in Fig. 2(b), D_1 snoops '0110' from the wavelength channels that are destined to D_2. The snooped data from D_1 can be transferred to a malicious core within the CMP to determine sensitive information. This type of snooping attack from malicious destination GIs is hard to detect, as it does not disrupt the intended communication among CMP cores. Therefore, there is a pressing need to address the security risks imposed by snooping GIs in DWDM-based PNoC architectures. To address this need, we propose a novel framework *SOTERIA* that improves hardware security in DWDM-based PNoC architectures.

4. *SOTERIA* FRAMEWORK: OVERVIEW

Our proposed multi-layer *SOTERIA* framework enables secure communication in DWDM-based PNoC architectures by integrating circuit-level and architecture-level enhancements. Fig. 3 gives a high-level overview of this framework. The PV-based security enhancement (*PVSC*) scheme uses the PV profile of the destination GIs' detector MRs to encrypt data before it is transmitted via the photonic waveguide. This scheme is sufficient to protect data from snooping GIs, if they do not know about the target destination GI. With target destination GI information, however, a snooping GI can decipher the encrypted data. Many PNoC architectures (e.g., [11], [27]) use the same waveguide to transmit both the destination GI information and actual data, making them vulnerable to data snooping attacks despite using *PVSC*. To further enhance security for these PNoCs, we devise an architecture-level reservation-assisted security enhancement (*RVSC*) scheme that uses a secure reservation waveguide to avoid the stealing of destination GI information by snooping GIs. The next two sections present details of our *PVSC* and *RVSC* schemes.

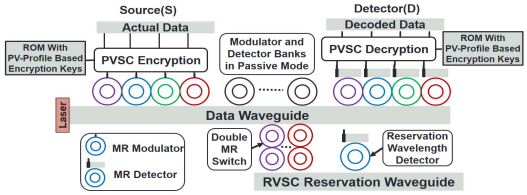

Fig. 3: Overview of proposed *SOTERIA* framework that integrates a circuit-level PV-based security enhancement (*PVSC*) scheme and an architecture-level reservation-assisted security enhancement (*RVSC*) scheme.

5. PV-BASED SECURITY ENHANCEMENT

As discussed earlier (Section 3.2), malicious destination GIs can snoop data from a shared waveguide. One way of addressing this security concern is to use data encryption so that the malicious destination GIs cannot decipher the snooped data. For the encrypted data to be truly undecipherable, the encryption key used for data encryption should be kept secret from the snooping GIs, which can be challenging as the identity of the snooping GIs in a PNoC is not known. Therefore, it becomes very difficult to decide whether or not to share the encryption key with a destination GI (that can be malicious) for data decryption. This conundrum can be resolved using a different key for every destination GI so that a key that is specific to a secure destination GI does not need to be shared with a malicious destination GI for decryption purpose. Moreover, to keep these destination specific keys secure, the malicious GIs in a PNoC must not be able to clone the algorithm (or method) used to generate these keys.

To generate unclonable encryption keys, our PV-based security (*PVSC*) scheme uses the PV profiles of the destination GIs' detector MRs. As discussed in [14], PV induces random shifts in the resonance wavelengths of the MRs used in a PNoC. These resonance shifts can be in the range from -3nm to 3nm [14]. The MRs that belong to different GIs in a PNoC have different PV profiles. In fact, the MRs that belong to different MR banks of the same GI also have different PV profiles. Due to their random nature, these MR PV profiles cannot be cloned by the malicious GIs, which makes the encryption keys generated using these PV profiles truly unclonable. Using the PV profiles of detector MRs, *PVSC* can generate a unique encryption key for each detector bank of every MWMR waveguide in a PNoC.

Our *PVSC* scheme generates encryption keys during the testing phase of the CMP chip, by using a dithering signal based in-situ method [15] to generate an anti-symmetric analog error signal for each detector MR of every detector bank that is proportional to the PV-induced resonance shift in the detector MR. Then, it converts the analog error signal into a 64-bit digital signal. Thus, a 64-bit digital error signal is generated for every detector MR of each detector bank. We consider 64 DWDM wavelengths per waveguide, and hence, we have 64 detector MRs in every detector bank and 64 modulator MRs in every modulator bank. For each detector bank, our *PVSC* scheme XORs the 64 digital error signals (of 64 bits each) from each of the 64 detector MRs to create a unique 64-bit encryption key. Note that our PVSC scheme also uses the same anti-symmetric error signals to control the carrier injection and heating of the MRs to remedy the PV-induced shifts in their resonances.

To understand how the 64-bit encryption key is utilized to encrypt data in photonic links, consider Fig. 4 which depicts an example photonic link that has one MWMR waveguide and connects the modulator banks of two source GIs (S_1 and S_2) with the detector banks of two destination GIs (D_1 and D_2). As there are two destination GIs on this link, *PVSC* creates two 64-bit encryption keys corresponding to them, and stores them at the source GIs. When data is to be transmitted by a source GI, the key for the appropriate destination is used to encrypt data at the flit-level granularity, by performing an XOR between the key and the data flit. This requires that the size of an encryption key match the data flit size. We consider the size of data flits to be 512 bits. Therefore, the 64-bit encryption key is appended eight times to generate a 512-bit encryption key. In Fig. 4, every source GI stores two 512-bit encryption keys (for destination GIs D_1 and D_2) in its local ROM, whereas every destination GI stores only its corresponding 512-bit key in its ROM. Note that we store the 512-bit keys instead of the 64-bit keys as this eliminates the latency overhead of affixing 64-bit keys to generate 512-bit keys, at the cost of a reasonable area/energy overhead in the ROM. As an example, if S_1 wants to send a data flit to D_2, then S_1 first accesses the 512-bit encryption key corresponding to D_2 from its local ROM and XORs the data flit with this key in one cycle, and then transmits the encrypted data flit over the link. As the link employs only one waveguide with 64 DWDM wavelengths, therefore, the encrypted 512-bit data flit is transferred on the link to D_2 in eight cycles. At D_2, the data flit is decrypted by XORing it with the 512-bit key corresponding to D_2 from the local ROM. In this scheme, even if D_1 snoops the data intended for D_2, it cannot decipher the data as it does not have access to the correct key (corresponding to D_2) for decryption. Thus, our *PVSC* encryption scheme protects data against snooping attacks in DWDM-based PNoCs.

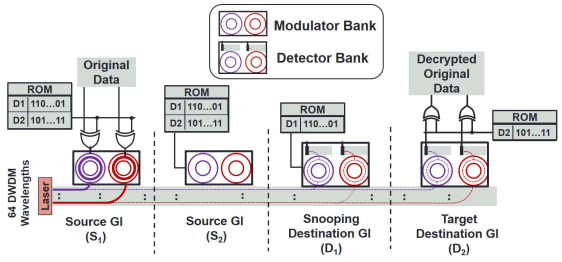

Fig. 4: Overview of proposed PV-based security enhancement scheme

Limitations of PVSC: The *PVSC* scheme can protect data from being deciphered by a snooping GI, if the following two conditions about the underlying PNoC architecture hold true: *(i)* the snooping GI does not know the target destination GI for the snooped data, *(ii)* the snooping GI cannot access the encryption key corresponding to the target destination GI. As discussed earlier, an encryption key is stored only at all source GIs and at the corresponding destination GI, which makes it physically inaccessible to a snooping destination GI.

However, if more than one GIs in a PNoC are compromised due to HTs in their control units and if these HTs launch a coordinated snooping attack, then it may be possible for the snooping GI to access the encryption key corresponding to the target destination GI.

For instance, consider the photonic link in Fig. 4. If both S_1 and D_1 are compromised, then the HT in S_1's control unit can access the encryption keys corresponding to both D_1 and D_2 from its ROM and transfer them to a malicious core (a core running a malicious program). Moreover, the HT in D_1's control unit can snoop the data intended for D_2 and transfer it to the malicious core. Thus, the malicious core may have access to the snooped data as well as the encryption keys stored at the source GIs. Nevertheless, accessing the encryption keys stored at the source GIs is not sufficient for the malicious GI (or core) to decipher the snooped data. This is because the compromised ROM typically has multiple encryption keys corresponding to multiple destination GIs, and choosing a correct key that can decipher data requires the knowledge of the target destination GI. Thus, our *PVSC* encryption scheme can secure data communication in PNoCs as long as the malicious GIs (or cores) do not know the target destinations of the snooped data.

Unfortunately, many PNoC architectures, e.g., [11], [27], that employ photonic links with multiple destination GIs utilize the same waveguide to transmit both the target destination information and actual data. In these PNoCs, if a malicious GI manages to tap the target destination information from the shared waveguide, then it can access the correct encryption key from the compromised ROM to decipher the snooped data. Thus, there is a need to conceal the target destination information from malicious GIs (cores). This motivates us to propose an architecture-level solution, as discussed next.

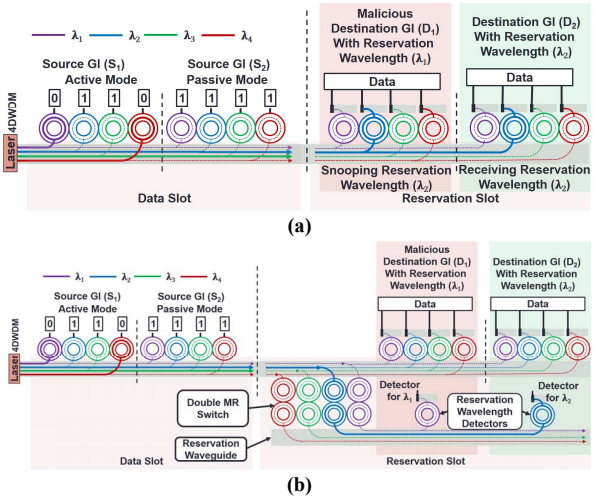

Fig. 5: Reservation-assisted data transmission in DWDM-based photonic waveguides (a) without *RVSC*, (b) with *RVSC*.

6. RESERVATION-ASSISTED SECURITY ENHANCEMENT

In PNoCs that use photonic links with multiple destination GIs, data is typically transferred in two time-division-multiplexed (TDM) slots called reservation slot and data slot [11], [27]. To minimize photonic hardware, PNoCs use the same waveguide to transfer both slots, as shown in Fig. 5(a). To enable reservation of the waveguide, each destination is assigned a reservation selection wavelength. In Fig. 5(a), $\boldsymbol{\lambda}_1$ and $\boldsymbol{\lambda}_2$ are the reservation selection wavelengths corresponding to destination GIs D_1 and D_2, respectively. Ideally, when a destination GI detects its reservation selection wavelength in the reservation slot, it switches ON its detector bank to receive data in the next data slot. But in the presence of an HT, a malicious GI can snoop signals from the reservation slot using the same detector bank that is used for data reception. For example, in Fig. 5(a), malicious GI D_1 is using one of its detectors to snoop $\boldsymbol{\lambda}_2$ from the reservation slot. By snooping $\boldsymbol{\lambda}_2$, D_1 can identify that the data it will snoop in the subsequent data slot will be intended for destination D_2. Thus, D_1 can now choose the correct encryption key from the compromised ROM to decipher its snooped data.

To address this security risk, we propose an architecture-level reservation-assisted security enhancement (*RVSC*) scheme. In *RVSC*, we add a reservation waveguide, whose main function is to carry reservation slots, whereas the data waveguide carries data slots. We use double MRs to switch the signals of reservation slots from the data waveguide to the reservation waveguide, as shown in Fig. 5(b). Double MRs are used instead of single MRs for switching to ensure that the switched signals do not reverse their propagation direction after switching [29]. Compared to single MRs, double MRs also have lower signal loss due to steeper roll-off of their filter responses [29].

The double MRs are switched ON only when the photonic link is in a reservation slot, otherwise they are switched OFF to let the signals of the data slot pass by in the data waveguide. Furthermore, in *RVSC*, each destination GI has only one detector on the reservation waveguide, which corresponds to its receiver selection wavelength. For example, in Fig. 5(b), D_1 and D_2 will have detectors corresponding to their reservation selection wavelengths $\boldsymbol{\lambda}_1$ and $\boldsymbol{\lambda}_2$, respectively, on the reservation waveguide. This makes it difficult for the malicious GI D_1 to snoop $\boldsymbol{\lambda}_2$ from the reservation slot as shown in Fig. 5(b), as D_1 does not have a detector corresponding to $\boldsymbol{\lambda}_2$ on the reservation waveguide. However, the HT in D_1's control unit may still attempt to snoop other reservation wavelengths (e.g., $\boldsymbol{\lambda}_2$) in the reservation slot by retuning D_1's $\boldsymbol{\lambda}_1$ detector. But succeeding in these attempts would require the HT to perfect the timing and target wavelength of its snooping attack, which is very difficult due to the large number of utilized reservation wavelengths. Thus, D_1 cannot identify the correct encryption key to decipher the snooped data. *In summary,* RVSC enhances security in PNoCs by protecting data from snooping attacks, even if the encryption keys used to secure data are compromised.

7. IMPLEMENTING *SOTERIA* FRAMEWORK ON PNOCS

We characterize the impact of *SOTERIA* on two popular PNoC architectures: Firefly [8] and Flexishare [9], both of which use DWDM-based photonic waveguides for data communication. We consider Firefly PNoC with 8×8 SWMR crossbar [8] and a Flexishare PNoC with 32×32 MWMR crossbar [9] with 2-pass token stream arbitration. We adapt the analytical equations from [29] to model the signal power loss and required laser power in the *SOTERIA*-enhanced Firefly and Flexishare PNoCs. At each source and destination GI of the *SOTERIA*-enhanced Firefly and Flexishare PNoCs, XOR gates are required to enable parallel encryption and decryption of 512-bit data flits. We consider a 1 cycle delay overhead for encryption and decryption of every data flit. The overall laser power and delay overheads for both PNoCs are quantified in the results section.

Firefly PNoC: Firefly PNoC [8], for a 256-core system, has 8 clusters (C1-C8) with 32 cores in each cluster. Firefly uses reservation-assisted SWMR data channels in its 8x8 crossbar for inter-cluster communication. Each data channel consists of 8 SWMR waveguides, with 64 DWDM wavelengths in each waveguide. To integrate *SOTERIA* with Firefly PNoC, we added a reservation waveguide to every SWMR channel. This reservation waveguide has 7 detector MRs to detect reservation selection wavelengths corresponding to 7 destination GIs. Furthermore, 64 double MRs (corresponding to 64 DWDM wavelengths) are used at each reservation waveguide to implement *RVSC*. To enable *PVSC*, each source GI has a ROM with seven entries of 512 bits each to store seven 512-bit encryption keys corresponding to seven destination GIs. In addition, each destination GI requires a 512-bit ROM to store its own encryption key.

Flexishare PNoC: We also integrate *SOTERIA* with the Flexishare PNoC architecture [9] with 256 cores. We considered a 64-radix 64-cluster Flexishare PNoC with four cores in each cluster and 32 data channels for inter-cluster communication. Each data channel has four

MWMR waveguides with each having 64 DWDM wavelengths. In *SOTERIA*-enhanced Flexishare, we added a reservation waveguide to each MWMR channel. Each reservation waveguide has 16 detector MRs to detect reservation selection wavelengths corresponding to 16 destination GIs. To enable *PVSC*, each source GI requires a ROM with 16 entries of 512 bits each to store the encryption keys, whereas each destination GI requires a 512-bit ROM.

8. EVALUATIONS

8.1. Evaluation Setup

To evaluate our proposed *SOTERIA* (*PVSC+RVSC*) security enhancement framework for DWDM-based PNoCs, we integrate it with the Firefly [8] and Flexishare [9] PNoCs, as explained in Section 7. We modeled and performed simulation based analysis of the *SOTERIA*-enhanced Firefly and Flexishare PNoCs using a cycle-accurate SystemC based NoC simulator, for a 256-core single-chip architecture at 22nm. We validated the simulator in terms of power dissipation and energy consumption based on results obtained from the DSENT tool [22]. We used real-world traffic from the PARSEC benchmark suite [23]. GEM5 full-system simulation [24] of parallelized PARSEC applications was used to generate traces that were fed into our 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 applications. We performed geometric calculations for a 20mm×20mm chip size, to determine lengths of SWMR and MWMR waveguides in Firefly and Flexishare. Based on this analysis, we estimated the time needed for light to travel from the first to the last node as 8 cycles at 5 GHz clock frequency [13]. We use a 512-bit packet size, as advocated in the Firefly and Flexishare PNoCs. Similar to [29], we adapt the VARIUS tool [20] to model random and systematic die-to-die (D2D) as well as within-die (WID) process variations in MRs for the Firefly and Flexishare PNoCs.

The static and dynamic energy consumption values for electrical routers and concentrators in Firefly and Flexishare PNoCs are based on results from DSENT [22]. We model and consider the area, power, and performance overheads for our framework implemented with the Firefly and Flexishare PNoCs as follows. *SOTERIA* with Firefly and Flexishare PNoCs has an electrical area overhead of 12.7mm^2 and 3.4mm^2, respectively, and power overhead of 0.44W and 0.36W, respectively, using gate-level analysis and CACTI 6.5 [25] tool for memory and buffers. The photonic area of Firefly and Flexishare PNoCs is 19.83mm^2 and 5.2mm^2, respectively, based on the physical dimensions [21] of their waveguides, MRs, and splitters. For energy consumption of photonic devices, we adapt model parameters from recent work [26], [28] with 0.42pJ/bit for every modulation and detection event and 0.18pJ/bit for the tuning circuits of modulators and photodetectors. The MR trimming power is 130µW/nm [30] for current injection and tuning power is 240µW/nm [30] for heating.

8.2. Overhead Analysis of *SOTERIA* on PNoCs

Our first set of experiments compare the baseline (without any security enhancements) Firefly and Flexishare PNoCs with their *SOTERIA* enhanced variants. From Section 7, all 8 SWMR waveguide groups of the Firefly PNoC and all 32 MWMR waveguide groups of the Flexishare PNoC are equipped with *PVSC* encryption/decryption and reservation waveguides for the *RVSC* scheme.

We adapt the analytical models from [29] to calculate the total signal loss at the detectors of the worst-case power loss node (N_{WCPL}), which corresponds to router C4R0 for the Firefly PNoC [8] and node R_{63} for the Flexishare PNoC [9]. Fig. 6(a) summarizes the worst-case signal loss results for the baseline and *SOTERIA* configurations for the two PNoC architectures. From the figure, Firefly PNoC with *SOTERIA* increases loss by 1.6dB and Flexishare PNoC with *SOTERIA* increases loss by 1.2dB on average, compared to their respective baselines. Compared to the baseline PNoCs that have no single or double MRs to switch the signals of the reservation slots, the double MRs used in the *SOTERIA*-enhanced PNoCs to switch the wavelength signals of the reservation slots increase through losses in the waveguides, which ultimately increases the worst-case signal losses in the *SOTERIA*-enhanced PNoCs. Using the worst-case signal losses shown in Fig. 6(a), we determine the total photonic laser power and corresponding electrical laser power for the baseline and *SOTERIA*-enhanced variants of Firefly and Flexishare PNoCs, shown in Fig. 6(b). From this figure, the Firefly and Flexishare PNoCs with *SOTERIA* have laser power overheads of 44.7% and 31.40% on average, compared to their baselines.

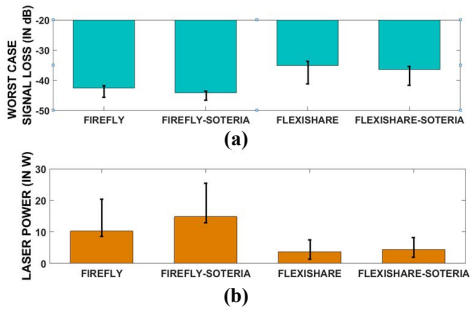

Fig. 6: Comparison of (a) worst-case signal loss and (b) laser power dissipation of *SOTERIA* framework on Firefly and Flexishare PNoCs with their respective baselines considering 100 process variation maps.

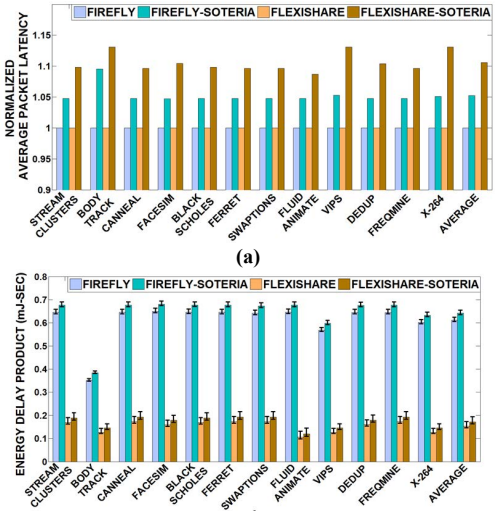

Fig. 7: (a) normalized average latency and (b) energy-delay product (EDP) comparison between different variants of Firefly and Flexishare PNoCs that include their baselines and their variant with *SOTERIA* framework, for PARSEC benchmarks. Latency results are normalized with their respective baseline 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.

Fig. 7 presents detailed simulation results that quantify the average packet latency and energy-delay product (EDP) for the two configurations of the Firefly and Flexishare PNoCs. Results are shown for twelve multi-threaded PARSEC benchmarks. From Fig. 7(a), Firefly with *SOTERIA* has 5.2% and Flexishare with *SOTERIA* has 10.6% higher latency on average compared to their respective baselines. The additional delay due to encryption and decryption of data (Section 7.1) with *PVSC* contributes to the increase in average latency.

From the results for EDP shown in Fig. 7(b), Firefly with *SOTERIA* has 4.9% and Flexishare with *SOTERIA* has 13.3% higher EDP on average compared to their respective baselines. Increase in EDP for the *SOTERIA*-enhanced PNoCs is not only due to the increase in their average packet latency, but also due to the presence of additional *RVSC* reservation waveguides, which increases the required photonic

hardware (e.g., more number of MRs) in the *SOTERIA*-enhanced PNoCs. This in turn increases static energy consumption (i.e., laser energy and trimming/tuning energy), ultimately increasing the EDP.

8.3. Analysis of Overhead Sensitivity

Our last set of evaluations explore how the overhead of *SOTERIA* changes with varying levels of security in the network. Typically, in a manycore system, only a certain portion of the data that contains sensitive information (i.e., keys) and only a certain number of communication links need to be secure. Therefore, for our analysis in this section, instead of securing all data channels of the Flexishare PNoC, we secure only a certain number channels using *SOTERIA*. Out of the total 32 MWMR channels in the Flexishare PNoC, we secure 4 (FLEX-ST-4), 8 (FLEX-ST-8), 16 (FLEX-ST-16), and 24 (FLEX-ST-24) channels, and evaluate the average packet latency and EDP for these variants of the *SOTERIA*-enhanced Flexishare PNoC.

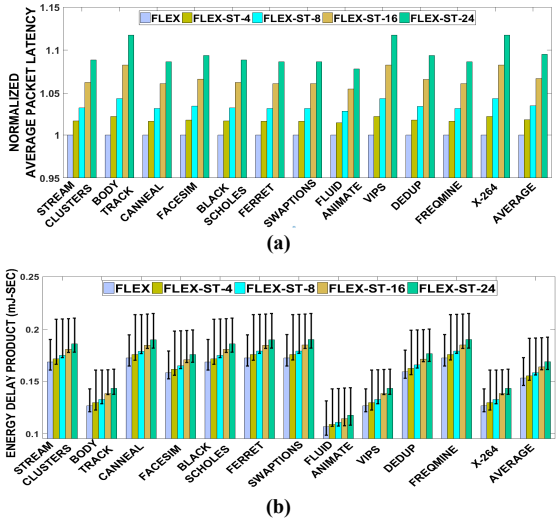

Fig. 8: (a) normalized latency and (b) energy-delay product (EDP) comparison between Flexishare baseline and Flexishare with 4, 8, 16, and 24 *SOTERIA* enhanced MWMR waveguide groups, for PARSEC benchmarks. Latency results are normalized to the baseline Flexishare results.

In Fig. 8, we present average packet latency and EDP values for the five *SOTERIA*-enhanced configurations of the Flexishare PNoC. From Fig. 8(a), FLEX-ST-4, FLEX-ST-8, FLEX-ST-16, and FLEX-ST-24 have 1.8%, 3.5%, 6.7%, and 9.5% higher latency on average compared to the baseline Flexishare. Increase in number of *SOTERIA* enhanced MWMR waveguides increases number of packets that are transferred through the *PVSC* encryption scheme, which contributes to the increase in average packet latency across these variants. From the results for EDP shown in Fig. 8(b), FLEX-ST-4, FLEX-ST-8, FLEX-ST-16, and FLEX-ST-24 have 2%, 4%, 7.6%, and 10.8% higher EDP on average compared to the baseline Flexishare. EDP in Flexishare PNoC increases with increase in number of *SOTERIA* enhanced MWMR waveguides. Increase in average packet latency and signal loss due to the higher number of reservation waveguides and double MRs increase overall EDP across these variants.

9. CONCLUSION

We presented a novel security enhancement framework called *SOTERIA* that secures data during unicast communications in DWDM-based PNoC architectures from snooping attacks. Our proposed *SOTERIA* framework shows interesting trade-offs between security, performance, and energy overhead for the Firefly and Flexishare PNoC architectures. Our analysis shows that *SOTERIA* enables hardware security in crossbar based PNoCs with minimal overheads of up to 10.6% in average latency and of up to 13.3% in EDP compared to the baseline PNoCs. Thus, *SOTERIA* represents an attractive solution to enhance hardware security in emerging DWDM-based PNoCs.

10. ACKNOWLEDGMENTS

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

REFERENCES

[1] R. Chakraborty, S. Narasimhan, S. Bhunia, "Hardware Trojan: Threats and emerging solutions," in Proc. HLDVT, pp. 166-171, Nov. 2009.
[2] M. Tehranipoor and F. Koushanfar, "A Survey of Hardware Trojan Taxonomy and Detection," IEEE Design & Test, pp. 10-25, Feb. 2009.
[3] S. Skorobogatov and C. Woods, "Breakthrough silicon scanning discovers backdoor in military chip," in Proc. CHES, pp. 23-40, Sept. 2012.
[4] W. J. Dally, B. Towles, "Route packets, not wires," in Proc. DAC, 2001.
[5] D. A. B. Miller, "Device requirements for optical interconnects to silicon chips," in JPROC, 97(7), pp. 1166-1185, 2009.
[6] C. Batten et al., "Building manycore processor-to-dram networks with monolithic silicon photonics," in HotI, pp. 21–30, 2008.
[7] I. Thakkar, S. V. R. Chittamuru, S. Pasricha, "Improving the Reliability and Energy-Efficiency of High-Bandwidth Photonic NoC Architectures with Multilevel Signaling," in Proc. NOCS, Oct. 2017.
[8] Y. Pan et al., "Firefly: Illuminating future network-on-chip with nanophotonics," in Proc. ISCA, 2009.
[9] Y. Pan, J. Kim, G.Memik, "Flexishare: Channel sharing for an energy efficient nanophotonic crossbar," in Proc. HPCA, 2010.
[10] S. V. R. Chittamuru, S. Pasricha, "SPECTRA: A Framework for Thermal Reliability Management in Silicon-Photonic Networks-on-Chip", in Proc. VLSID, Jan 2016.
[11] D. M. Ancajas, et al., "Fort-NoCs: Mitigating the Threat of a Compromised NoC," in Proc. DAC, 2014.
[12] C. Li, et al., "Energy-efficient optical broadcast for nanophotonic networks-on-chip," in Proc. OIC, pp. 64-65, 2012.
[13] S. V. R. Chittamuru, S. Desai, S. Pasricha, "SWIFTNoC: A reconfigurable silicon photonic network with multicast enabled channel sharing for multicore architectures," in ACM JETC, 13(4), no. 58, 2017.
[14] S. K. Selvaraja., "Wafer-Scale Fabrication Technology for Silicon Photonic Integrated Circuits," PhD thesis, Ghent University, 2011.
[15] K. Padmaraju et al., "Wavelength Locking and Thermally Stabilizing Microring Resonators Using Dithering Signals," in JLT, 32 (3), 2013.
[16] C. H. Gebotys, et al., "A framework for security on NoC technologies," in Proc. ISVLSI, Feb. 2003.
[17] H. K. Kapoor, et al., "A Security Framework for NoC Using Authenticated Encryption and Session Keys," in CSSP, 2013.
[18] S. V. R. Chittamuru, I. Thakkar, S. Pasricha, "Process Variation Aware Crosstalk Mitigation for DWDM based Photonic NoC Architectures," in Proc. ISQED, Mar. 2016.
[19] S. V. R. Chittamuru, I. Thakkar, S. Pasricha, "PICO: Mitigating Heterodyne Crosstalk Due to Process Variations and Intermodulation Effects in Photonic NoCs," in Proc. DAC, Jun. 2016.
[20] S. Sarangi et al., "Varius: A model of process variation and resulting timing errors for microarchitects," IEEE TSM, 21(1), pp. 3 –13, 2008.
[21] S. Xiao, M. H. Khan, H. Shen, and M. Qi, "Modeling and measurement of losses in silicon-on-insulator resonators and bends," in Optics Express, 15(17), pp. 10553-10561, 2007.
[22] C. Sun et al., "DSENT - a tool connecting emerging photonics with electronics for opto-electronic networks-on-chip modeling," NOCS, 2012.
[23] C. Bienia et al., "The PARSEC Benchmark Suit: Characterization and Architectural Implications," in PACT, Oct. 2008.
[24] N. Binkert et al., "The gem5 Simulator," in CA News, May 2011.
[25] CACTI 6.5, http://www.hpl.hp.com/research/cacti/
[26] S. V. R. Chittamuru, I. Thakkar, S. Pasricha, "Analyzing Voltage Bias and Temperature Induced Aging Effects in Photonic Interconnects for Manycore Computing," in Proc. SLIP, June. 2017.
[27] C. Chen and A. Joshi, "Runtime management of laser power in silicon-photonic multibus NoC architecture," in Proc. IEEE JQE, 2013.
[28] I. Thakkar, S. V. R. Chittamuru, S. Pasricha, "Mitigation of Homodyne Crosstalk Noise in Silicon Photonic NoC Architectures with Tunable Decoupling," in Proc. CODES+ISSS, Oct. 2016.
[29] S. V. R. Chittamuru, I. Thakkar, S. Pasricha, "HYDRA: Heterodyne Crosstalk Mitigation with Double Microring Resonators and Data Encoding for Photonic NoCs," in TVLSI, vo1. 26, no. 1, 2018.
[30] D. Dang, S. V. R. Chittamuru, R. Mahapatra, and S. Pasricha, "Islands of Heaters: A Novel Thermal Management Framework for Photonic NoCs," in Proc. ASPDAC, Jan. 2017.