Integration, the VLSI Journal 68 (2019) 22–29

Contents lists available at ScienceDirect

Integration, the VLSI Journal

journal homepage: www.elsevier.com/locate/vlsi

A thermal-sensitive design of a 3D torus-based optical NoC architecture☆

Yaoyao Ye *, Zhe Zhang

Department of Micro/Nano Electronics, School of Electronic Information and Electrical Engineering, Shanghai Jiao Tong University, Shanghai, 200240, PR China

A R T I C L E I N F O

Keywords:
Chip multiprocessor
Optical network-on-chip
Thermal effect
Routing algorithm

A B S T R A C T

In order to overcome limitations of traditional electronic interconnects in terms of power efficiency and bandwidth density, optical networks-on-chip (NoC) based on silicon photonics have been proposed as an emerging on-chip communication architecture for chip multiprocessors (CMPs) with large core counts. However, due to thermo-optic effects, wavelength-selective silicon photonic devices such as microresonators suffer from temperature-dependent wavelength shift. In this work, we propose a thermal-sensitive design of a 3D torus-based optical NoC architecture. For the 3D torus-based optical NoC architecture, we propose a hybrid optical-electronic router architecture with a fully-connected 7×7 optical switching fabric. Besides, a thermal-sensitive routing algorithm is proposed to optimize the optical power loss in the presence of on-chip temperature variations. Simulation results show that for a set of synthetic traffic patterns, the proposed 3D torus-based optical NoC with the thermal-sensitive routing reduces the average thermal-induced optical power loss by 14.3% and 17% respectively as compared to the matched 3D torus-based optical NoC and the 3D mesh-based optical NoC with the traditional XYZ routing. For a set of real applications, the proposed 3D torus-based optical NoC with the thermal-sensitive routing reduces the worst-case optical power loss by 7.9% and 14.6% respectively as compared to the matched 3D torus-based optical NoC and the 3D mesh-based optical NoC with the traditional XYZ routing.

1. Introduction

As the scale of transistors enters the deep nanometer region, the number of transistors available on a single chip has increased to several billions. By enabling energy-efficient parallel processing at lower clock frequencies, chip multiprocessors (CMPs) are a natural platform for embedded systems as well as high-performance computing. Network-on-chip (NoC) architectures have been widely proposed as a new generation of on-chip communication architectures, which could scale better than on-chip shared buses and ad-hoc networks as the number of cores increases [1]. However, due to the limitations of traditional electronic interconnects in power efficiency and bandwidth density, as well as issues of high-frequency crosstalk noise and parasitic capacitance in deep-submicron integrated circuit design, there are still bandwidth, power efficiency and reliability bottlenecks in traditional NoCs based on electronic interconnects.

With the booming developments in nanoscale silicon photonic technologies for short-haul communications, silicon photonics based optical interconnects are emerging as a promising new approach to moving on-chip data at high speeds and low power. As compared to the traditional electronic interconnects, optical interconnects can enable significantly increased bandwidth density, low power consumption, and low latency. By integrating optical interconnects in NoC architectures, optical NoCs can overcome many of the most serious on-chip communication issues [2–5]. Most of the prior works on optical NoCs are based on silicon photonic devices including optical waveguides and silicon microresonators (MRs). Considering the high efficiency of electronic interconnects in on-chip local communication as well as for control, optical NoCs are often controlled and configured in the electronic domain. Three-dimensional (3D) integration technologies provide the support for realizing mixed-technology electronic-controlled optical NoCs.

However, one of the major challenges in optical NoC designs is thermal sensitivity, which is an intrinsic characteristic of photonic devices. Due to the fact that the power density on chip is uneven and the thermal conductivity of packaging materials is limited, chip temperature fluctuates temporally as well as spatially. The temperature can rise quickly from room temperature after a cold start, and vary by more than 30 °C across a steady-state chip under typical operating conditions [6]. As a result of thermo-optic effects, wavelength-selective silicon photonic devices such as microresonators, which are widely used in optical NoCs,

☆ This work was supported by the National Natural Science Foundation of China under Grant 61602298, and was sponsored by Shanghai Sailing Program.
* Corresponding author.
E-mail address: yeyaoyao@sjtu.edu.cn (Y. Ye).

https://doi.org/10.1016/j.vlsi.2019.05.007
Received 15 February 2019; Received in revised form 22 April 2019; Accepted 27 May 2019
Available online 1 June 2019
0167-9260/

suffer from temperature-dependent wavelength shifts [7]. An investigation of related thermal issues shows that if we take the thermal regulation power into account, optical interconnects may not have advantages in power efficiency as compared with their electrical counterparts [8].

In order to tackle the thermal issue in optical NoC designs, in this work, we propose a thermal-sensitive design of a 3D torus-based optical NoC architecture in the presence of on-chip temperature variations. For the 3D torus-based optical NoC architecture, we propose a hybrid optical-electronic router architecture with a fully-connected 7×7 optical switching fabric. Besides, we propose a thermal-sensitive routing algorithm to optimize its optical power loss in the presence of on-chip temperature variations. The rest of the paper is organized as follows. In Section 2, we review the thermal sensitivity in optical NoCs. In Section 3, we present the proposed thermal-sensitive design and power optimization methodology. Section 5 presents simulation results and comparisons in terms of thermal-induced energy efficiency and network performance. Last, Section 6 concludes this work.

2. Related work

In the last decade, optical NoCs have been proposed as a new approach to empowering ultra-high bandwidth with low power consumption [9]. Kirman et al. proposed a hierarchical optical bus for multiprocessor systems [10]. The optical loop encircles the chip with WDM (wavelength division multiplexing) support at the top level of the hierarchy. Shacham et al. proposed an augmented-torus optical NoC based on a blocking 4×4 optical switch [11]. The network is circuit-switched and some critical network design issues such as path-setup and tear-down are covered in the work. Pasricha et al. proposed to use an optical ring waveguide with bus protocol standards to replace global pipelined electrical interconnects [12]. Pan et al. proposed an optical crossbar-based on-chip network architecture with localized arbitrations [13]. It supports high throughput with multiple global crossbars. Cianchetti et al. proposed an optical network with a pre-decoded source routing mechanism [14]. Ding et al. presented an optical routing framework to reduce the power consumption of ONoCs [15]. Batten et al. proposed an optical mesh based on a hybrid optical-electrical global crossbar, where the processing cores and DRAM are divided into sub-meshes and connected with the optical crossbar [16]. Gu et al. proposed an optical NoC architecture named MRONoC [17]. By employing sufficient wavelengths resources and making an efficient wavelength arrangement, MRONoC can support contention-free communications and simplify the arbitration. The authors also proposed an optical NoC architecture with time-division multiplexing and wavelength-division multiplexing technology to solve the network contention problem. The number of wavelength groups and timeslots is optimized by using a genetic algorithm [18]. Cerutti et al. proposed a scheduler for multi-wavelength ring-based optical NoCs [19]. An iterative parallel wavelength matching algorithm is implemented in the scheduler to address the wavelength assignment problem.

Previous works have investigated the thermal sensitivity of photonic devices used in optical NoCs. The temperature-dependent wavelength shift in silicon MR is found to be non-negligible in practical use [7]. In order to mitigate thermal effects in optical NoCs, run-time thermal management techniques such as OS-based workload migration and DVFS (dynamic voltage and frequency scaling) have been proposed to reduce the on-chip temperature gradients [20–22]. Due to the limitations of these thermal management techniques, device-level thermal compensation techniques are still in need. Thermal tuning by local microheaters is one such device-level solution, however, it is relatively slow and power inefficient [23]. Additionally, microresonators with low temperature dependence as well as athermal microresonators have been demonstrated by applying proper polymer materials [24]. However, there are still compatibility issues when implementing athermal microresonators with CMOS technology. Furthermore, some efforts have been made to overcome the thermal challenges in optical NoCs from system-level perspectives [25–28]. In Ref. [25], the authors proposed a thermal-aware methodology to design optical NoCs with distributed CMOS-compatible VCSELs, based on steady state thermal simulations and SNR (signal-to-noise ratio) analysis. In Refs. [26,27], the authors systematically modeled thermal effects in optical NoCs and proposed several low temperature-sensitivity techniques. In Ref. [28], a system-level proactive thread migration technique and a device-level thermal island framework were proposed to alleviate the thermal issues in optical NoCs.

In order to further alleviate the thermal issue in optical NoCs, in this work, we propose a thermal-sensitive design of a 3D torus-based optical NoC architecture in the presence of on-chip temperature variations. This paper is an extended version of work published in Ref. [29]. In Ref. [29], we proposed a 3D torus-based optical NoC architecture with a thermal-sensitive source routing, where the whole path is determined at the source according to global temperature information. We extend our previous work [29] by a distributed thermal-sensitive adaptive routing.

3. A 3D torus-based optical NoC architecture

In this section, we propose a thermal-sensitive design approach for a 3D torus-based optical NoC architecture (Fig. 1). Based on optical thermal models that characterize thermal effects in optical NoCs from a system-level perspective, a thermal-sensitive routing mechanism is proposed for the 3D torus-based optical NoC to optimize the optical power loss in the presence of on-chip temperature variations.

3.1. Network architecture

Regular network topologies, such as mesh and torus, are preferred in NoC designs because of their predictable scalability in terms of performance and power consumption. In order to improve the NoCs performance with shorter interconnection lengths, NoC architectures based on 3D network topologies have been proposed [30–32]. It has been demonstrated that as compared to its 2D implementation, the 3D mesh-based NoC can improve performance significantly with higher integration densities and smaller footprints. As compared to 3D mesh topology, the 3D torus topology takes advantage of the wrap-around links among edge nodes to offer better path diversity and better load balance.

Fig. 2 shows the topology of a 3D torus-based optical NoC architecture, where each router is connected to a local processor core. Each processor is assigned a unique ID of for addressing, and the local router has the same address. For many-core chip multiprocessors, a hierarchical hybrid optical-electrical NoC architecture would be a good solution by taking advantages of electrical interconnect for short-distance traffic. In this work, we focus on solving the problem of thermal issues in optical NoCs. The proposed thermal-sensitive design of 3D torus-based optical NoC actually can also be used as the optical part of hierarchical hybrid optical-electrical NoC architectures. The proposed 3D torus-based optical

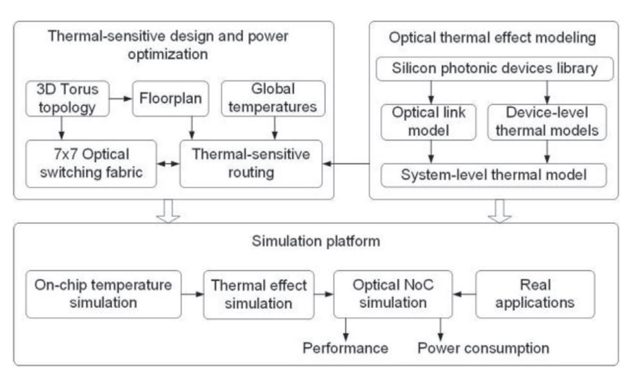

Fig. 1. The proposed thermal-sensitive design approach for a 3D torus-based optical NoC architecture.

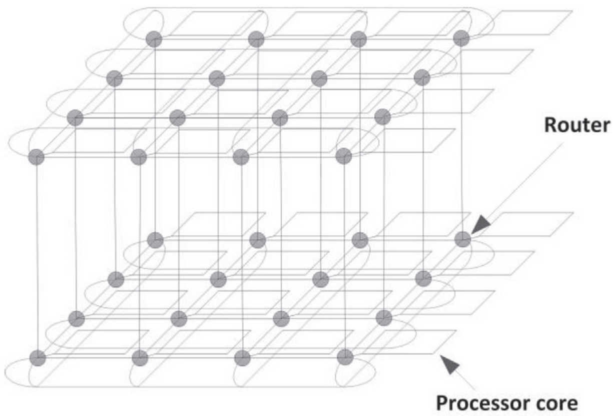

Fig. 2. A 3D torus-based optical NoC architecture.

NoC uses circuit switching, in which an optical path is reserved before payload transmission. An overlapped electronic control network is used for optical path configuration and maintenance. Before a packet transmission, a single-flit path-setup packet would be routed in the electronic control network for path reservation. The payload data is transmitted along the reserved optical path after the path setup. High-speed optical transmission is achieved in this architecture without buffering in intermediate routers.

3.2. Hybrid optical-electronic router architecture

In the 3D torus-based optical NoC, hybrid optical-electronic routers are used to interconnect processor cores. The hybrid optical-electronic router architecture (Fig. 3) is composed of a 7×7 fully-connected optical switch, an electronic control unit, and an electronic/optical (E/O) interface. The electronic control unit includes a thermal-sensitive routing unit and an adaptive power control unit. The 7×7 fully-connected optical switch is responsible for directing optical signals in the network, while the electronic control unit implements the routing algorithm and power control. In addition, the electronic control units in the network are interconnected into an electronic control network with metallic interconnects for control purposes, which includes selecting optical path configurations and delivering temperature information. The E/O interface, which is in charge of serialization, deserialization, and E/O conversions, is used to facilitate communications between the electronic and optical domain.

In Ref. [32], a reduced 7×7 optical switch was proposed specially for dimension-order routing in a 3D mesh-based optical NoC. In this paper, we extend it to a 7×7 fully-connected optical switch (Fig. 5), which can support any routing algorithms including the newly proposed thermal sensitive routing scheme. Compared to the reduced 7×7 optical switch for dimension-order routing only proposed in Ref. [32], our new port-to-port switching functions in the proposed 7×7 fully-connected

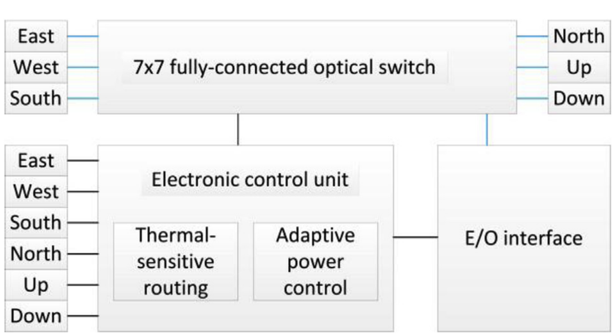

Fig. 3. A hybrid optical-electronic router architecture.

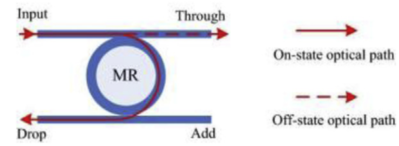

Fig. 4. A 1×2 basic optical switching fabric.

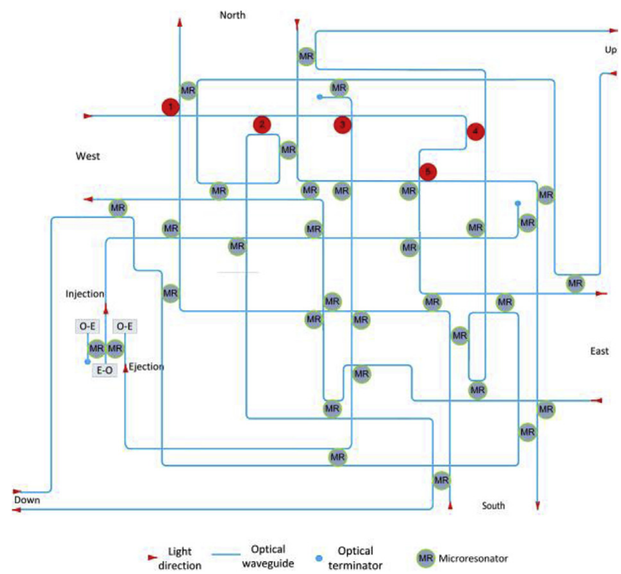

Fig. 5. A 7×7 fully-connected optical switching fabric.

optical switching fabric include switchings from south to east/west, from north to east/west, from up to east/west/south/north, from down to east/west/south/north.

The 7×7 fully-connected optical switch has seven bidirectional ports, including injection/ejection, up, down, east, west, south, and north ports. The ports are aligned to their intended directions. The local processor core is connected by the injection/ejection ports through an O/E interface while other ports are connected to neighboring optical switches. The 7×7 fully-connected optical switch is built from basic optical switching elements which are based on microresonators, optical waveguides, and optical terminators [32]. Fig. 4 shows the 1×2 basic optical switching fabric. By powering on/off the microresonator, the basic optical switching element implements 1×2 optical switching functions. When the microresonator is in the on state, input signal would be coupled into the ring and be directed to the drop port. When the microresonator is in the off state, input signal would propagate along the input waveguide to the through port. In Fig. 5, we mark the key microresonators in red for switchings from the west to other ports. In detail, by turning on the microrsonator 1/2/3/4/5, input signal from the west port can be switched to the north/down/ejection/up/south respectively. The proposed 7×7 fully-connected optical switch inherits the passive routing feature from the reduced 7×7 optical switch [32]. For switching from the west to the east output port, no microresonators need to be turned on.

In order to get the router-level optical loss model from the device-level optical loss model, we propose to model the 7×7 optical switching fabric by three 7×7 matrixes, A, P, and C. For an optical signal switched from the input port i to the output port j, the elements A_{ij}, P_{ij} and C_{ij} represent the number of active switching, the number of passive switching, and the number of waveguide crossing respectively. By using the new defined three matrixes, the router-level optical loss from the input port i to the output port j could be expressed as Equation (1), where we assume the microresonators in the same router are with the same temperature. L_{active} is the optical power loss inserted in an active switching, $L_{passive}$ is the optical power loss inserted in a passive switching, and L_{cro} is the optical power loss inserted by a waveguide crossing.

$$L_{ij} = A_{ij} \cdot L_{active} + P_{ij} \cdot L_{passive} + C_{ij} \cdot L_{cro} \tag{1}$$

4. A thermal-sensitive routing mechanism

For mesh or torus-based NoCs, the traditional dimension-order routing is a low-complexity deterministic routing algorithm. However, due to the deterministic feature, the traditional dimension-order routing cannot avoid bad paths with significant thermal-induced optical power loss in the presence of on-chip temperature variations. In this section, we propose a thermal-sensitive routing algorithm to find optimal paths with the minimum estimated thermal-induced optical power loss according to runtime on-chip temperature distributions.

4.1. Thermal-sensitive optical power loss

Due to the thermo-optic effect, material refractive index is temperature dependent and follows Equation (2), where n_0 is the refractive index at room temperature, dn/dT is the thermo-optic coefficient of the material, and ΔT is the temperature variation. Physical measurements show that the thermo-optic coefficient of silicon is on the order of $10^{-4}/K$ and is nonlinear over a large temperature range at 1550 nm wavelength [33].

$$\mathrm{n} = n_0 + \Delta T \cdot \frac{dn}{dT} \tag{2}$$

As a result of thermo-optic effect, wavelength-selective silicon photonic devices such as microresonators, which are widely used in optical NoCs, suffer from temperature dependent wavelength shifts [7]. The emission wavelength of on-chip lasers, such as VCSELs (vertical cavity surface-emitting lasers), also shifts with ambient temperatures [34], while the power efficiency degrades at high temperatures [35]. The thermal related wavelength mismatch between the laser located in the source node and the microresonators in intermediate nodes along a photonic path can cause significant additional optical power loss. Besides, the output power of a VCSEL degrades at higher operating temperatures. If using an off-chip VCSEL as the laser source, the lasing wavelength can be fixed by equipping with a temperature control unit. If using an on-chip VCSEL as the laser source, the temperature-dependent wavelength shift and power efficiency degradation should both be taken into account in the thermal model.

4.2. Thermal-sensitive routing mechanism

The routing algorithm determines a path to transmit a packet from the source to the destination. In traditional deterministic routing such as dimension-order routing, a deterministic path is decided based only on the source address and the destination address. However, in the presence of on-chip temperature variations, some paths which are under high temperature variations suffer from severe optical power loss. In order to avoid such paths, we propose to introduce adaptiveness in routing by considering on-chip temperature conditions when making routing decisions.

In the earlier conference version of this work [29], we proposed a thermal-sensitive source routing, where the whole path is determined at the source. Global temperature information is required to make routing decisions, so the overhead includes the broadcasting of global temperature information. A temperature table keeps the global temperature information at the granularity of each processor core. Considering that the on-chip temperature could fluctuate temporally, the temperature table will receive updates from other nodes. The routing unit includes a shortest path routing unit, a temperature table, and an optical thermal-effect modeling unit. For each packet to be transmitted, the shortest path routing unit reads the header flit of the packet to get the source id and destination id, and then it finds out the shortest paths between the source and destination. The optical thermal-effect modeling unit gets all the candidate paths information from the shortest path routing unit, and reads the temperature table for global temperature information. It then calculates the thermal-induced optical power loss for each candidate path. All the shortest paths are considered as candidate paths, and they will be compared for thermal-induced optical power loss by the optical thermal-effect modeling unit. The path with the minimum thermal-induced optical power loss will be selected out of all the candidate paths.

In this work, we propose a distributed thermal-sensitive routing mechanism for loss reduction in the proposed 3D torus-based optical NoC. The proposed thermal-sensitive routing is based on a learning-based adaptive routing algorithm proposed in Ref. [36]. We assume the optical NoC is circuit switching, in which an optical path is reserved before payload transmission. For each packet to be sent, a setup packet is routed in the control network. In traditional Q-routing based electronic NoCs, Q-values represent the latency of alternative paths, and the Q-learning technique is performed to learn network congestions for latency optimization [37]. In this work, we use the Q-learning technique to learn the runtime on-chip temperature information and find an optimal shortest path in a distributed way in order to reduce optical power loss. We target optical NoCs in presence of temperature variations, and focus on the optimization of thermal-induced optical power loss. In the proposed routing, each router in the optical NoC keeps an L-table which stores L-values. The L-values represent the estimated thermal-induced optical power loss of alternative paths. Each router learns the state of the network by receiving and storing L-values from neighboring routers.

Assume that a setup packet is generated at the source node s, and its destination is node d. Assume that this setup packet just arrives at an intermediate router y, and the previous router is x. The receiver at the router y will store it in buffer. The router controller is responsible for routing decision and arbitration. The router controller at the router y first gets a set of valid output ports along shortest paths, which guarantees that alternative paths found are shortest and the proposed routing can achieve the power optimization at little sacrifice of network performance. Then the optimal selection function selects an optimal output port with the minimum estimated thermal-induced optical power loss towards the destination node. After the current router forwards the setup packet to the next router through the selected output port, it sends back a learning packet which contains the L-value to the previous router. Upon receiving the learning packet, the previous router updates the corresponding L-value in local L-table. Assume that N(y) is the set of feasible next routers found along the shortest path, we define $L_y(\mathrm{n}, \mathrm{d})$ as the L-value representing the estimated thermal-induced optical power loss from any feasible next router $\mathrm{n} \varepsilon \mathrm{N}(\mathrm{y})$ to the destination. The neighboring router $\mathrm{z} \varepsilon \mathrm{N}(\mathrm{y})$ which is with the minimum estimated thermal-induced optical power loss will be selected as the next router.

If the setup packet in the buffer has been granted with an outport, it will be popped to the output line and ready for transmission. The sender at the local router will then check the output line in every clock cycle, and pass the setup packet to the next router. As soon as the current router forwards the setup packet to the next router, it will send a L-value back to the previous router, which indicates the minimum estimated thermal-induced optical power loss from the current router to the destination. Upon receiving the estimated minimum L-value, the previous router will compute and update a new L-value according to Equation (3), where γ is the learning rate. The learning rate determines the convergence speed of the proposed routing. The learning rate is set as 1 in our simulations. It will update the local L table with the new L-value.

$$L_x(y, d)_{new} = L_x(y, d)_{old} + \gamma \cdot (L_x(y, d)_{min} - L_x(y, d)_{old}) \tag{3}$$

5. Simulation results and comparisons

We developed a SystemC-based cycle-accurate network simulator for a 3D $8 \times 8 \times 2$ torus-based optical NoC with the proposed thermal-sensitive routing algorithm. Network simulations were conducted under several synthetic traffic patterns as well as a set of real applications. For each real application, we used traffic information to simulate the on-chip temperature distributions with McPAT [38] and HotSpot [39]. We assumed 40 Gbps data-link bandwidth for the optical NoC. The electronic control network operates at 1.25 GHz with 32-bit wide bidirectional metallic interconnects. In order to show the efficiency of the proposed thermal-sensitive routing algorithm, we used a matched 3D $8 \times 8 \times 2$ torus-based optical NoC as well as a 3D $8 \times 8 \times 2$ mesh-based optical NoC, both with the traditional XYZ routing as baselines for comparisons.

5.1. The convergence of the proposed routing algorithm

In order to show the convergence of the proposed thermal-sensitive routing algorithm, we simulated a 3D $8 \times 8 \times 2$ torus-based optical NoC with the proposed routing algorithm under uniform traffic. A random temperature distribution was applied during the simulation, and we assumed that the on-chip temperature was within the range of [55 °C, 85 °C]. We assumed the 3-dB bandwidth of the basic optical switching element is 1.24 nm. Fig. 6 shows the convergence of thermal-induced optical power loss optimization for delivering packets from node 0 to node 91, from node 0 to node 82, and from node 0 to node 73. It is shown that for packets belongs to the same source-destination communication pair, the thermal-induced optical power loss in path varies before the proposed routing algorithm reaches the state of convergence. After learning the network with a certain number of packets, the proposed routing algorithm reaches the state of convergence, which means that it finds the optimal path with optimized thermal-induced optical power loss. For packets sent from node 0 to node 91, it finds the optimal path after learning five packets, and the optimal path is with optical power loss of 26.9 dB. For packets sent from node 0 to node 73, it finds an optimal path with optical power loss of about 19.5 dB. If we assume a smaller 3 dB bandwidth, the optimization of optical power loss would be more significant. It can be observed from Fig. 6 that it needs to learn less packets to find the optimal path for a shorter path. Fig. 7 shows another example convergence procedure of the proposed routing under corner block temperature distribution, where more than half of the hot cores are located at the corners. It shows that the proposed routing algorithm can also converge quickly.

5.2. Thermal-induced optical power loss and energy efficiency

In the following, we present more simulation results and comparisons in terms of thermal-induced optical power loss and energy efficiency, under several synthetic traffic patterns as well as a set of real applications.

The power consumed in optical domain includes the power consumed

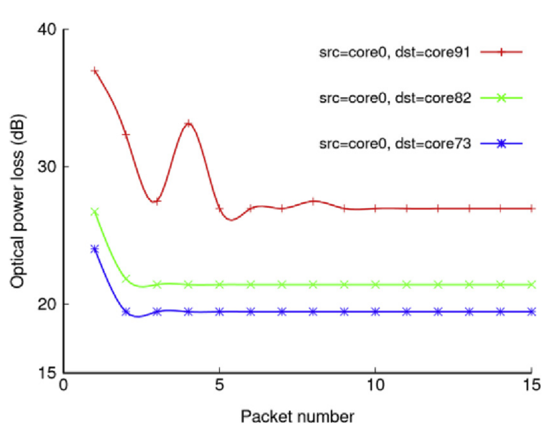

Fig. 6. The convergence of thermal-induced optical power loss optimization by the proposed routing, random temperature distribution.

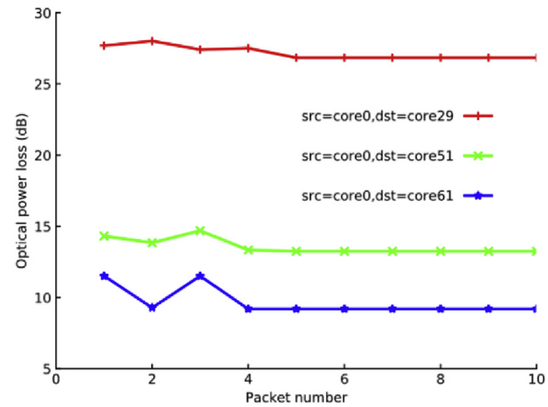

Fig. 7. The convergence of thermal-induced optical power loss optimization by the proposed routing, corner block temperature distribution.

by turning on microresonators in optical routers and the power consumed by optical/electronic (O/E) interfaces for O/E conversions. We assume the power consumption for turning on a microresonator is 20 μW [40]. For evaluation of the thermal-induced optical power loss inserted by a microresonator, we assume the 3 dB bandwidth is 1.24 nm. A typical O/E interface includes serializer, driver, VCSEL, waveguide, photodetector, TIALA circuits and deserializer. Energy consumption for EO and OE conversions in an optical link is the sum of power consumed by all components of the O/E interface. In optical NoCs, power dissipated in the O/E interface is mainly governed by the laser source. We assume to use VCSEL as the off-chip laser source [35]. We assume that the VCSELs are directly modulated by driving currents and no external modulation is needed during optical transmission. We assume that if the driving current I is above the threshold current of 2.5 mA, output power will increase approximately linearly with the driving current with slope efficiency of 0.36 mW/mA. VCSEL power consumption can be calculated as UI, where U is the bias voltage and is also assumed to increase linearly with the driving current. According to the 10 Gbps driver and TIA-LA circuits demonstrated in Ref. [41], the energy efficiency of driver is scaled from 0.2pJ/bit in 80 nm to 0.1125pJ/bit in 45 nm, and the energy efficiency of TIA-LA circuits is scaled from 0.6pJ/bit in 80 nm to 0.3375pJ/bit in 45 nm. The energy efficiency of serializer and deserializer is scaled from 0.576pJ/bit in 90 nm to 0.288pJ/bit in 45 nm [42]. The photodetector model is based on a 10 Gbps Ge waveguide photodetector monolithically integrated in 130 nm CMOS process with a sensitivity of −14.2 dBm for 10^{-12} of bit error rate (BER) [43]. To ensure that optical NoCs function properly, the optical signal power received by the receiver at the destination should be equal or greater than the receiver sensitivity. This condition must hold, otherwise the bit error rate would increase. We assume to use an adaptive power control mechanism, which sets the driving current of the laser source according to the loss in the path.

5.2.1. Synthetic traffic patterns

We simulated a 3D $8 \times 8 \times 2$ torus-based optical NoC with the proposed thermal-sensitive routing algorithm under a set of synthetic traffic patterns. A matched 3D $8 \times 8 \times 2$ torus-based optical NoC as well as a 3D $8 \times 8 \times 2$ mesh-based optical NoC both with the traditional XYZ routing were simulated as a baseline for comparisons. During the simulations, we applied random on-chip temperature distributions within the range of [55 °C, 85 °C].

Fig. 8 shows the comparison of the average thermal-induced optical power loss of all packets under each synthetic traffic pattern. The injection rate is set as 0.1. It is shown that on average of the four synthetic traffic patterns, the 3D torus-based optical NoC with the proposed thermal-sensitive routing reduces the average thermal-induced optical power loss by 14.3% and 17% respectively as compared to the matched 3D torus-based optical NoC and the 3D mesh-based optical NoC with the

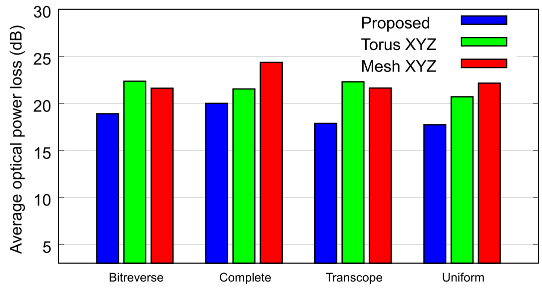

Fig. 8. The average thermal-induced optical power loss under synthetic traffic patterns, with random temperature distributions.

traditional XYZ routing. If we assume a smaller 3-dB bandwidth, the improvement of optical power loss would be more significant. It could also be observed that the proposed routing algorithm has a greater optimization space for traffic patterns with more long-distance traffic (e.g. Bitreverse) than the uniform traffic.

Fig. 9 shows the comparison of the normalized average thermal-induced optical energy efficiency of all packets under each synthetic traffic pattern. The injection rate is set as 0.1. It is shown that the proposed 3D torus-based optical NoC with the thermal-sensitive routing reduces the average thermal-induced optical energy consumption significantly. On average of the four synthetic traffic patterns, it reduces the average thermal-induced optical energy consumption by 76.1% and 84.8% respectively as compared to the matched 3D torus-based optical NoC and the 3D mesh-based optical NoC with the traditional XYZ routing.

5.2.2. *Real applications*

In addition to the synthetic traffic patterns, we also simulated the 3D $8 \times 8 \times 2$ torus-based optical NoC with the proposed 3D Torus-based optical NoC with the thermal sensitive routing under a set of real applications. We used the NoC traffic suit [44] to simulate real applications. Specifically, an offline optimization approach was applied for each real application to map and schedule tasks onto the multiprocessor system-on-chip with the objective of maximizing the system performance. Our network simulated with the generated traffic traces. A matched 3D $8 \times 8 \times 2$ torus-based optical NoC as well as a 3D $8 \times 8 \times 2$ mesh-based optical NoC with the traditional XYZ routing were simulated as a baseline for comparisons. For each real application, we used traffic information to simulate the on-chip temperature distribution with HotSpot [39] and McPAT [38].

Fig. 10 shows the comparisons of the worst-case thermal-induced optical power loss among all packets in each real application. The worst-case thermal-induced optical power loss for the proposed method is defined as the worst-case thermal-induced optical power loss among all the packet transmissions after optimization. It shows that on average of the five real applications, the proposed 3D torus-based optical NoC with the thermal-sensitive routing reduces the worst-case optical power loss by 7.9% and 14.6% respectively as compared to the matched 3D torus-based optical NoC and the 3D mesh-based optical NoC with the traditional XYZ routing. For applications with more long-distance traffic (e.g. the SATELL application), the proposed routing algorithm has a greater space for optimization, and thus it could achieve more significant improvements. This can be further illustrated by Fig. 11, which shows the loss distribution of the SATELL application. It is shown that the loss distribution is shifted to the low-loss regions by using the proposed thermal-sensitive routing. With the traditional XYZ routing, 10.1% of packets are with more than 30 dB loss. If using the proposed routing, only 1.7% of packets are with more than 30 dB loss. Fig. 12 also shows the same trend.

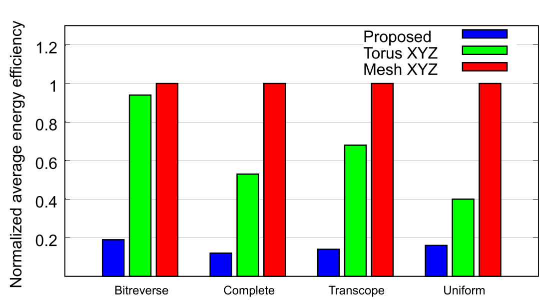

Fig. 9. The average thermal-induced optical energy efficiency under synthetic traffic patterns.

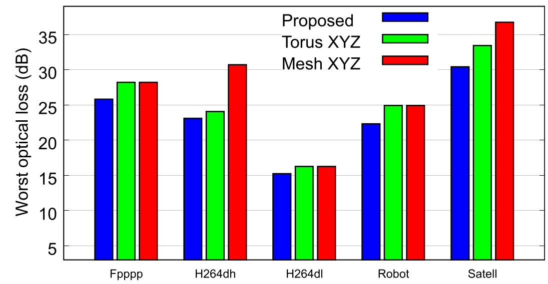

Fig. 10. The worst-case thermal-induced optical power loss under real applications, with simulated temperature distributions.

Fig. 13 shows the comparisons of the normalized worst-case thermal-induced optical energy efficiency among all packets in each real application. The worst-case energy efficiency for the proposed method is defined as the worst-case energy efficiency among all the packet transmissions after optimization. It shows that on average of the five real applications, the proposed 3D torus-based optical NoC with the thermal-sensitive routing reduces the worst-case energy consumption by 43.8% and 62.8% respectively as compared to the matched 3D torus-based optical NoC and the 3D mesh-based optical NoC with the traditional XYZ routing.

Fig. 14 shows that on average of the five real applications, the proposed 3D torus-based optical NoC with the thermal-sensitive routing reduces the average optical power loss by 3.7% and 5.6% respectively as

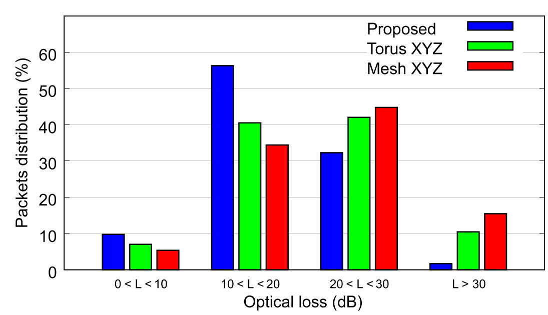

Fig. 11. The loss distribution of the SATELL application.

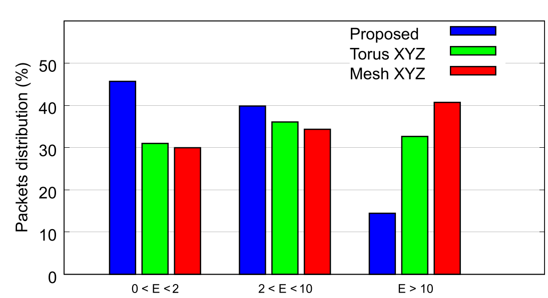

Fig. 12. The power distribution of the SATELL application.

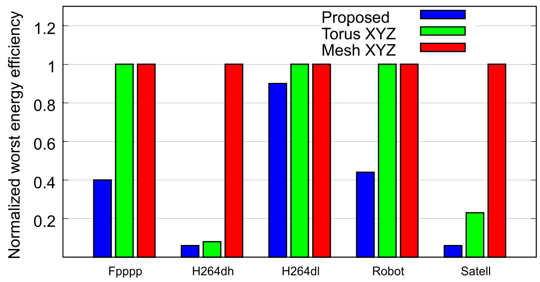

Fig. 13. The worst-case thermal-induced optical energy efficiency under real applications.

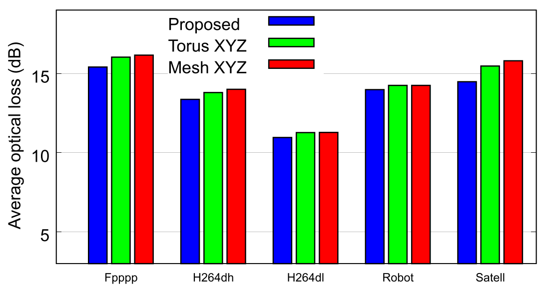

Fig. 14. The average thermal-induced optical loss under real applications, with simulated temperature distributions.

compared to the matched 3D torus-based optical NoC and the 3D mesh-based optical NoC with the traditional XYZ routing. Correspondingly, Fig. 15 shows that on average of the five real applications, the proposed 3D torus-based optical NoC with the thermal-sensitive routing reduces the average optical energy efficiency by 8.2% and 14.6% respectively as compared to the matched 3D torus-based optical NoC and the 3D mesh-based optical NoC with the traditional XYZ routing. For applications with more long-distance traffic, e.g. the SATELL application, the proposed routing reduces the average thermal-induced optical energy consumption by 27.2% as compared to the XYZ routing. We could conclude that the proposed routing has a greater space for optimization for long-distance traffic as compared to short-distance traffic. If we assume a smaller 3-dB bandwidth, the improvement of optical power loss would be more significant.

Since the proposed thermal-sensitive routing finds the optimal path from shortest paths, it achieves the power optimization at little sacrifice of network performance. This can be illustrated by Fig. 16, which shows the comparisons of normalized network performance under each real application. Here the network performance is measured in terms of the application end-to-end (ETE) delay. It is shown that the proposed routing optimizes the optical loss at the cost of some overhead on network performance. For some applications, the proposed thermal-sensitive routing would lead to a larger end-to-end delay than the traditional XYZ routing.

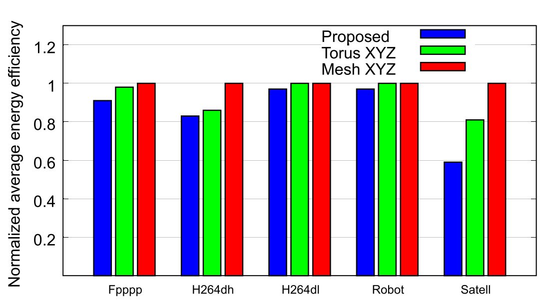

Fig. 15. The average thermal-induced optical energy efficiency under real applications.

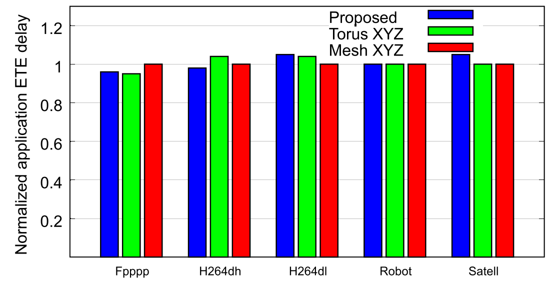

Fig. 16. The normalized performance under real applications.

On average of the five real applications, the application ETE delay increases by 0.25% as compared to the traditional XYZ routing. In this work, we focus on reducing the optical power loss. In future work, we may extend the current work to a two-fold optimization problem for both the optical power loss and the communication delay.

6. Conclusions

In order to tackle the thermal issues in optical NoC designs, we propose a thermal-sensitive design of a 3D torus-based optical NoC architecture. For the 3D torus-based optical NoC architecture, we propose a hybrid optical-electronic router architecture with a fully-connected 7×7 optical switching fabric. Besides, a thermal-sensitive routing algorithm is proposed to optimize its power consumption in the presence of on-chip temperature variations. Simulation results show that for a set of four synthetic traffic patterns, the proposed 3D torus-based optical NoC with the thermal-sensitive routing reduces the average thermal-induced optical power loss by 14.3% and 17% respectively as compared to the matched 3D torus-based optical NoC and the 3D mesh-based optical NoC with the traditional XYZ routing. For a set of five real applications, the proposed 3D torus-based optical NoC with the thermal-sensitive routing reduces the worst-case optical power loss by 7.9% and 14.6% respectively as compared to the matched 3D torus-based optical NoC and the 3D mesh-based optical NoC with the traditional XYZ routing. It could also be concluded that the proposed routing algorithm has a greater optimization space for traffic patterns with more long-distance traffic. The proposed thermal-sensitive design of 3D torus-based optical NoC can also be used as the optical part of hierarchical hybrid optical-electrical NoC architectures.

References

[1] L. Benini, G. De Micheli, Networks on chip: a new paradigm for systems on chip design, in: Proc. Design, Automation and Test in Europe Conference and Exhibition, 2002, pp. 418–419.
[2] S. Bahirat, S. Pasricha, UC-PHOTON: a novel hybrid photonic network-on-chip for multiple use-case applications, in: 11th Int. Symp. Quality Electronic Design, 2010, pp. 721–729.
[3] L. Ramini, A. Ghiribaldi, P. Grani, S. Bartolini, H.T. Fankem, D. Bertozzi, Assessing the energy break-even point between an optical noc architecture and an aggressive electronic baseline, in: 2014 Design, Automation Test in Europe Conference Exhibition (DATE), March 2014, pp. 1–6.
[4] S. Werner, J. Navaridas, M. Lujan, Amon: an advanced mesh like optical noc, in: 2015 IEEE 23rd Annual Symposium on High Performance Interconnects, Aug 2015, pp. 52–59.
[5] S. Faralli, N. Andriolli, F. Gambini, P. Pintus, G. Preve, M. Chiesa, R. Ortuno, O. Liboiron-Ladouceur, I. Cerutti, Bidirectional transmissions in a ring-based packaged optical noc with 12 add-drop microrings, in: 2016 IEEE Photonics Conference (IPC), Oct 2016, pp. 621–622.
[6] K. Skadron, M.R. Stan, K. Sankaranarayanan, W. Huang, S. Velusamy, D. Tarjan, Temperature-aware microarchitecture: modeling and implementation, ACM Trans. Archit. Code Optim. 1 (1) (2004) 94–125.
[7] V. Padgaonkar, A. Arbor, Thermal Effects in Silicon Based Resonant Cavity Devices, NNIN REU Research Accomplishments, 2004.
[8] R.K. Dokania, A.B. Apsel, Analysis of challenges for on-chip optical interconnects, in: Proc. 19th ACM Great Lakes Symp. VLSI, 2009, pp. 275–280.
[9] T. Alexoudi, N. Terzenidis, S. Pitris, M. Moralis-Pegios, C. Vagionas, C. Mitsolidou, G. Mourgias-Alexandris, G.T. Kanellos, A. Miliou, K. Vyrsokinos, N. Pleros, Optics in computing: from photonic network-on-chip to chip-to-chip

interconnects and disintegrated architectures, J. Light. Technol. 37 (2) (Jan 2019) 363–379.
[10] N. Kirman, M. Kirman, R. Dokania, J. Martinez, A. Apsel, M. Watkins, D. Albonesi, On-chip optical technology in future bus-based multicore designs, IEEE Micro 27 (1) (Jan.-Feb. 2007) 56–66.
[11] A. Shacham, K. Bergman, L. Carloni, On the design of a photonic network-on-chip, in: First Int. Symp. Networks-on-Chip, 2007, pp. 53–64.
[12] S. Pasricha, N. Dutt, ORB: an on-chip optical ring bus communication architecture for multi-processor systems-on-chip, in: Asia and South Pacific Conf. Design Automation, Mar. 2008, pp. 789–794.
[13] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, A. Choudhary, Firefly: illuminating future network-on-chip with nanophotonics, in: Proc. 36th Int. Symp. Computer Architecture (ISCA), 2009.
[14] M.J. Cianchetti, J.C. Kerekes, D.H. Albonesi, Phastlane: a rapid transit optical routing network, in: Proc. 36th Annual. Inter. Symp. Computer Architecture, 2009, pp. 441–450.
[15] D. Ding, Y. Zhang, H. Huang, R. Chen, D. Pan, O-router: an optical routing framework for low power on-chip silicon nano-photonic integration, in: 46th ACM/IEEE Design Automation Conf., July 2009, pp. 264–269.
[16] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. Holzwarth, M. Popovic, H. Li, H. Smith, J. Hoyt, F. Kartner, R. Ram, V. Stojanovic, K. Asanovic, Building manycore processor-to-DRAM networks with monolithic silicon photonics, in: 16th IEEE Symp. High Performance Interconnects, 2008, pp. 21–30.
[17] H. Gu, K. Chen, Y. Yang, Z. Chen, B. Zhang, MRONoC: a low latency and energy efficient on chip optical interconnect architecture, IEEE Photon. J. 9 (1) (Feb 2017) 1–12.
[18] H. Gu, Z. Wang, B. Zhang, Y. Yang, K. Wang, Time-divisionmultiplexing-wavelength-division-multiplexing-based architecture for ONoC, IEEE/OSA J. Opt. Commun. Netw. 9 (5) (May 2017) 351–363.
[19] I. Cerutti, M.N.A. Acmad, R. Reyes, P. Castoldi, N. Andriolli, Scheduling in multi-wavelength ring-based optical networks-on-chip, IEEE/OSA J. Opt. Commun. Netw. 10 (4) (April 2018) 322–331.
[20] Z. Li, M. Mohamed, X. Chen, E. Dudley, K. Meng, L. Shang, A.R. Mickelson, R. Joseph, M. Vachharajani, B. Schwartz, Y. Sun, Reliability modeling and management of nanophotonic on-chip networks, Very Large Scale Integr. (VLSI) Syst., IEEE Trans. PP (99) (2010) 1–14.
[21] T. Zhang, J.L. Abellan, A. Joshi, A.K. Coskun, Thermal management of manycore systems with silicon-photonic networks, in: 2014 Design, Automation Test in Europe Conference Exhibition (DATE), March 2014, pp. 1–6.
[22] S.V.R. Chittamuru, S. Pasricha, Spectra: a framework for thermal reliability management in silicon-photonic networks-on-chip, in: 2016 29th International Conference on VLSI Design and 2016 15th International Conference on Embedded Systems (VLSID), Jan 2016, pp. 86–91.
[23] F. Gan, T. Barwicz, M. Popovic, M. Dahlem, C. Holzwarth, P. Rakich, H. Smith, E. Ippen, F. Kartner, Maximizing the thermo-optic tuning range of silicon photonic structures, in: Photonics in Switching, Aug. 2007, pp. 67–68.
[24] V. Raghunathan, J. Hu, W.N. Ye, J. Michel, L.C. Kimerling, Athermal silicon ring resonators, in: Integrated Photonics Research, Silicon and Nanophotonics, 2010, p. IMC5.
[25] H. Li, A. Fourmigue, S.L. Beux, X. Letartre, I. O'Connor, G. Nicolescu, Thermal aware design method for vcsel-based on-chip optical interconnect, in: 2015 Design, Automation Test in Europe Conference Exhibition (DATE), March 2015, pp. 1120–1125.
[26] Y. Ye, J. Xu, X. Wu, W. Zhang, X. Wang, M. Nikdast, Z. Wang, W. Liu, System-level modeling and analysis of thermal effects in optical networks-on-chip, Very Large Scale Integr. (VLSI) Syst., IEEE Trans. PP (99) (2012) 1–14.
[27] Y. Ye, Z. Wang, P. Yang, J. Xu, X. Wu, X. Wang, M. Nikdast, Z. Wang, L.H.K. Duong, System-level modeling and analysis of thermal effects in WDM-based optical networks-on-chip, IEEE Trans. Comput. Aided Des. Integr Circuits Syst. 33 (11) (Nov 2014) 1718–1731.
[28] D. Dang, S.V.R. Chittamuru, R. Mahapatra, S. Pasricha, Islands of heaters: a novel thermal management framework for photonic NoCs, in: 2017 22nd Asia and South Pacific Design Automation Conference (ASP-DAC), Jan 2017, pp. 306–311.
[29] K. Yao, Y. Ye, S. Pasricha, J. Xu, Thermal-sensitive design and power optimization for a 3D torus-based optical NoC, in: 2017 IEEE/ACM International Conference on Computer-Aided Design (ICCAD), Irvine, CA, 2017, pp. 827–834.
[30] B. Feero, P. Pande, Networks-on-chip in a three-dimensional environment: a performance evaluation, Comput., IEEE Trans. 58 (1) (Jan. 2009) 32–45.
[31] H. Gu, J. Xu, Design of 3D optical network on chip, in: Photonics and Optoelectronics, Symposium on, Aug. 2009, pp. 1–4.
[32] Y. Ye, J. Xu, B. Huang, X. Wu, W. Zhang, X. Wang, M. Nikdast, Z. Wang, W. Liu, Z. Wang, 3-D mesh-based optical network-on-chip for multiprocessor system-on-chip, IEEE Trans. Comput. Aided Des. Integr Circuits Syst. 32 (4) (April 2013) 584–596.
[33] F.G. Della Corte, M. Esposito Montefusco, L. Moretti, I. Rendina, G. Cocorullo, Temperature dependence analysis of the thermos-optic effect in silicon by single and double oscillator models, J. Appl. Phys. 88 (12) (Dec. 2000) 7115–7119.
[34] S. Mogg, N. Chitica, U. Christiansson, R. Schatz, P. Sundgren, C. Asplund, M. Hammar, Temperature sensitivity of the threshold current of long-wavelength InGaAs-GaAs VCSELs with large gain cavity detuning, IEEE J. Quantum Electron. 40 (5) (May. 2004) 453–462.
[35] A. Syrbu, A. Mereuta, V. Iakovlev, A. Caliman, P. Royo, E. Kapon, 10GbpsVCSELswithhighsinglemodeoutputin1310nm and 1550 nm wavelength bands, in: Conf. Optical Fiber communication/National Fiber Optic Engineers, 2008, pp. 1–3.
[36] Z. Zhang, Y. Ye, A learning-based thermal-sensitive power optimization approach for optical NoCs, ACM J. Emerg. Technol. Comput. Syst. (JETC) 14 (2) (July 2018) no. 21.
[37] F. Farahnakian, M. Ebrahimi, M. Daneshtalab, P. Liljeberg, J. Plosila, Q-learning based congestion-aware routing algorithm for on-chip network, in: 2011 IEEE 2nd International Conference on Networked Embedded Systems for Enterprise Applications, Dec 2011, pp. 1–7.
[38] S. Li, J.H. Ahn, R.D. Strong, J.B. Brockman, D.M. Tullsen, N.P. Jouppi, McPat: an integrated power, area, and timing modeling framework for multicore and manycore architectures, in: 2009 42nd Annual IEEE/ACM International Symposium on Microarchitecture (MICRO), Dec 2009, pp. 469–480.
[39] W. Huang, S. Ghosh, S. Velusamy, K. Sankaranarayanan, K. Skadron, M.R. Stan, Hotspot: a compact thermal modeling methodology for early-stage VLSI design, IEEE Trans. Very Large Scale Integr. (VLSI) Syst. 14 (5) (May 2006) 501–513.
[40] A.W. Poon, F. Xu, X. Luo, Cascaded active silicon microresonator array cross-connect circuits for WDM networks-on-chip, Silicon Photon. III 6898 (1) (2008).
[41] C. Kromer, G. Sialm, C. Berger, T. Morf, M. Schmatz, F. Ellinger, D. Erni, G.-L. Bona, H. Jackel, A 100-mW 4 x 10 Gb/s transceiver in 80-nm CMOS for high-density optical interconnects, IEEE J. Solid State Circuits 40 (12) (Dec. 2005) 2667–2679.
[42] J. Poulton, R. Palmer, A. Fuller, T. Greer, J. Eyles, W. Dally, M. Horowitz, A 14-mW 6.25-Gb/s transceiver in 90-nm CMOS, IEEE J. Solid State Circuits 42 (12) (Dec. 2007) 2745–2757.
[43] G. Masini, G. Capellini, J. Witzens, C. Gunn, A 1550nm, 10Gbps monolithic optical receiver in 130nm CMOS with integrated Ge waveguide photodetector, in: 4th IEEE Int. Conf. Group IV Photonics, 2007, pp. 1–3.
[44] W. Liu, J. Xu, X. Wu, Y. Ye, X. Wang, W. Zhang, M. Nikdast, Z. Wang, A NoC traffic suite based on real applications,, in: IEEE Computer Society Annual Symp. VLSI, 2011.