3-D Mesh-Based Optical Network-on-Chip for Multiprocessor System-on-Chip

Yaoyao Ye, *Student Member, IEEE,* Jiang Xu, *Member, IEEE,* Baihan Huang, *Student Member, IEEE,* Xiaowen Wu, *Student Member, IEEE,* Wei Zhang, *Member, IEEE,* Xuan Wang, *Student Member, IEEE,* Mahdi Nikdast, *Student Member, IEEE,* Zhehui Wang, *Student Member, IEEE,* Weichen Liu, *Member, IEEE,* and Zhe Wang, *Student Member, IEEE*

***Abstract*—Optical networks-on-chip (ONoCs) are emerging communication architectures that can potentially offer ultrahigh communication bandwidth and low latency to multiprocessor systems-on-chip (MPSoCs). In addition to ONoC architectures, 3-D integrated technologies offer an opportunity to continue performance improvements with higher integration densities. In this paper, we present a 3-D mesh-based ONoC for MPSoCs, and new low-cost nonblocking 4 × 4, 5 × 5, 6 × 6, and 7 × 7 optical routers for dimension-order routing in the 3-D mesh-based ONoC. Besides, we propose an optimized floorplan for the 3-D mesh-based ONoC. The floorplan follows the regular 3-D mesh topology but implements all optical routers in a single optical layer. The floorplan is optimized to minimize the number of extra waveguide crossings caused when merging the 3-D ONoC to one optical layer. Based on a set of real applications and uniform traffic pattern, we develop a SystemC-based cycle-accurate NoC simulator and compare the 3-D mesh-based ONoC with the matched 2-D mesh-based ONoC and 2-D electronic NoC for performance and energy efficiency. Additionally, we quantitatively analyze thermal effects on the 3-D 8 × 8 × 2 mesh-based ONoC.**

***Index Terms*—3-D, floorplan, mesh, multiprocessor, optical network-on-chip, optical router.**

I. Introduction

AS THE SCALE of transistors enter the nanometer region, the number of transistors available on a single chip increases to billions or even greater. Multiprocessor system-on-chip (MPSoC) is becoming an attractive platform for high-performance applications. The performance of the multiprocessor systems is determined not only by the performance of individual processors, but also by how efficiently they collaborate with one another. An efficient on-chip communication architecture can help fully utilize the computation resources offered by multiple processors. The on-chip communication architectures have gradually moved to network-on-chip (NoC) to alleviate the problems of poor scalability, limited bandwidth, and high power consumption in traditional interconnection architectures [1], [2]. As semiconductor technologies continually shrink feature sizes, metallic interconnects gradually become the bottleneck of NoC performance due to the limited bandwidth, long delay, large area, and high power dissipation.

Optical interconnects have demonstrated their advantages in multicomputer systems and on-board inter-chip interconnections. With the booming developments in nanoscale silicon photonic technologies for short-haul communications, optical NoCs (ONoCs) are also proposed as an emerging communication architecture for the new-generation multiprocessor systems. They can potentially offer ultrahigh communication bandwidth, low latency, and high energy efficiency. Different ONoC architectures have been proposed based on optical waveguides and microresonators (MRs) [3]–[7].

In addition to NoC architectures, 3-D integrated technologies are another attractive solution for overcoming the barriers to interconnect scaling. Among several 3-D integration technologies, the through-silicon via (TSV) approach is the most promising one and therefore is the focus of the majority of 3-D integration research activities [8]. In a TSV-based 3-D chip, multiple device layers are stacked together with direct vertical interconnects tunneling through them. The 3-D integration technologies make possible the realization of mixed-technology electronic-controlled ONoCs. The electronic-controlled ONoCs were proposed to be implemented in a TSV-based two-layer 3-D chip, where an optical device layer is stacked upon the CMOS electronic layer through TSVs [9]. Based on the 3-D integration technologies, researchers also proposed some 3-D NoC architectures based on 3-D network topologies to improve the NoCs performance with shorter interconnection lengths [10], [11].

In this paper, we propose a 3-D mesh-based ONoC for MPSoCs, and new 6 × 6 and 7 × 7 optical routers for dimension-order routing in the 3-D mesh-based ONoC. New 4 × 4 and 5 × 5 routers are also proposed to be used in the network corners and edges. The routers are built based on silicon MRs and are strictly nonblocking. By fully utilizing the properties of dimension-order routing in the 3-D mesh,

Manuscript received February 20, 2012; revised July 27, 2012; accepted October 22, 2012. Date of current version March 15, 2013. This work was supported in part by RPC11EG18 and SBI06/07.EG01-4. This paper was recommended by Associate Editor Y. Xie.

Y. Ye, J. Xu, B. Huang, X. Wu, X. Wang, M. Nikdast, Z. Wang, W. Liu, and Z. Wang are with the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Kowloon, Hong Kong (e-mail: yeyaoyao@ust.hk; jiang.xu@ust.hk; bhuangaa@ust.hk; wxxaf@ust.hk; eexwang@ust.hk; mnikdast@ust.hk; zhehui@ust.hk; weichen@ust.hk; zwangag@ust.hk).

W. Zhang is with the School of Computer Engineering, Nanyang Technological University, Singapore 639798 (e-mail: zhangwei@ntu.edu.sg).

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

Digital Object Identifier 10.1109/TCAD.2012.2228739

the new routers minimize the number of MRs and reduce the number of waveguide crossings inside the switching fabrics. In addition, we propose optimized floorplans for the 3-D mesh-based ONoCs. The floorplan follows the regular 3-D mesh topology but implements all optical routers in a single optical device layer. The floorplan of the optical device layer is designed to minimize the number of extra waveguide crossings caused when merging the 3-D optical routing network to one optical layer. As an intrinsic characteristic of photonic devices, thermal sensitivity is a potential issue in ONoC designs. Based on the system-level analytical ONoC thermal model proposed in [12], we quantitatively study the thermal effects on the 3-D $8 \times 8 \times 2$ mesh-based ONoC based on the new proposed low-cost optical routers.

The rest of this paper is organized as follows. Section II gives a survey of related works on ONoCs. Section III details the design of new optical routers for dimension-order routing in the 3-D mesh-based ONoCs. The characteristics of the new routers are analyzed and compared with other optical routers. The floorplan design for the 3-D mesh-based ONoCs is presented in Section IV. Simulation results are then analyzed and compared in Section V. We compare the 3-D mesh-based ONoC with the matched 2-D mesh-based ONoC and the 2-D mesh-based electronic NoC under real applications and uniform traffic. We also quantitatively analyze the thermal effects on the 3-D $8 \times 8 \times 2$ mesh-based ONoC. Section VI concludes this paper.

II. Related Work

ONoC offers a new approach to empowering ultrahigh bandwidth with low power consumption. Kirman *et al.* [3] proposed a hierarchical optical bus for multiprocessor systems. The optical loop encircles the chip with wavelength division multiplexing support at the top level of the hierarchy. Briere *et al.* [13] presented a multistage ONoC with a passive-switching λ-router. Beausoleil *et al.* [14] proposed a crossbar-based ONoC, where 64 wavelengths are multiplexed over 270 waveguides. Two hundred fifty six waveguides are allocated for control and data, and 14 waveguides are for broadcast and arbitration. Shacham *et al.* [6] proposed an augmented-torus ONoC based on a blocking 4×4 optical switch. The network is circuit-switched and some critical network design issues such as path-setup and tear-down are covered in the work. Mo *et al.* [5] proposed a hierarchical mesh-based ONoC. It uses hybrid optical–electronic routers for electronic wormhole switching in local networks and circuit switching in the global mesh-based optical network. Pasricha *et al.* [15] proposed to use an optical ring waveguide with bus protocol standards to replace global pipelined electrical interconnects. Pan *et al.* [4] proposed an optical crossbar-based on-chip network architecture with localized arbitrations. This supports high throughput with multiple global crossbars. Gu *et al.* [16] proposed a fat-tree based ONoC. Cianchetti *et al.* [17] proposed an optical network with a predecoded source routing mechanism. Low latency is achieved by transmitting packets several hops in a single clock cycle if no contention exists. Kodi *et al.* [18] proposed an ONoC architecture for 64 cores. The bandwidth is maximized with the proposed static routing and wavelength allocation schemes. Ding *et al.* [19] presented an optical routing framework to reduce the power consumption of ONoCs. Batten *et al.* [7] proposed an optical mesh based on a hybrid optical–electrical global crossbar, where the processing cores and DRAM are divided into submeshes and connected with the optical crossbar.

For the on-chip optical routers, MRs are commonly used as a wavelength-selective optical switch to perform the switching functions. Silicon MRs of small size ($5\,\mu$m radius) have been demonstrated. The fabrication is based on silicon waveguides with a cross-section of 500×200 nm and the insertion loss is about 0.5 dB [20]. Experiment results show that the MR on–off switching power consumption is on the order of $20\,\mu$W [21]. Based on the switching functions of MRs, several 5×5 optical routers have been proposed for routings in mesh or torus-based ONoCs [22]–[25]. Cygnus is a low-power nonblocking 5×5 optical router [22]. It uses 16 MRs to implement a 5×5 switching function. The 5×5 nonblocking crossbar allows five concurrent transactions if there is no contention for the same output port. Cygnus does not need to turn on any MR if an optical signal travels in the same dimension through the router, such as from north to south, or from west to east. Only one MR will be powered on if the optical signal turns from one dimension to another or uses the injection/ejection port. Regardless of the network size, at most three MRs will be powered on in any XY-routing optical path in a 2-D optical mesh or torus, including one MR for injection at the source, one MR for a turn during the routing, and one MR for ejection at the destination node. Based on Cygnus, a more compact 5×5 nonblocking optical router Crux, was proposed [23]. It inherits the passive routing feature of Cygnus. But the switching functions of Crux are reduced to XY routing only in mesh or torus-based ONoCs. An optimized crossbar-based 5×5 optical router was demonstrated in [24]. One MR is required to be turned on in the optimized crossbar for every switching. Ji *et al.* [25] demonstrated another nonblocking 5×5 optical router. The router uses 16 MRs in total. One MR should be turned on for every switching, except for switchings from east to south, or south to east, or west to north, or north to west.

In order to combine the benefits of NoCs and the 3-D integration, 3-D electronic NoC architectures based on 3-D network topologies have been proposed in the literature. Feero *et al.* [10] evaluated the performance of 3-D electronic NoCs with 3-D mesh and fat-tree topologies. They show that as compared to their 2-D implementations, the 3-D realization of both mesh and fat-tree based NoCs could improve the performance significantly with higher integration densities and smaller footprints. Besides the mesh and tree-based topologies, other network topologies such as De Bruijn Graph and honeycomb are also alternatives. And these topologies have also been extended to the 3-D paradigm for NoC designs. Chen *et al.* [26] proposed a 3-D NoC architecture based on De Bruijn Graph to achieve smaller network diameter and smaller network latency. There are also some research works on the 3-D ONoC architectures. Gu *et al.* [11] proposed a 3-D optical cubic-mesh NoC together with an optimized partial crossbar-

based optical router. In this paper, we design new passive-routing 6×6 and 7×7 optical routers and propose optimized floorplans for the 3-D mesh-based ONoCs. New 4×4 and 5×5 optical routers are also designed to be used in the network edges. The number of waveguide crossings is minimized in router designs and floorplan optimizations.

Thermal sensitivity is an intrinsic characteristic of photonic devices used by ONoCs as well as a potential issue. As a result of thermo-optic effect, the temperature-dependent wavelength shifts in vertical cavity surface-emitting lasers (VCSELs) and silicon-based MRs are found to be about 50–100 pm/°C [27], [28]. As a widely used device in ONoCs, MRs perform as a wavelength-selective optical switch or modulator. The thermal related wavelength variations will result in additional optical power loss. Besides, VCSEL power efficiency degrades at high temperatures [29]. The work [12] systematically modeled and quantitatively analyzed the thermal effects in ONoCs. Several techniques were proposed to reduce the temperature sensitivity of ONoCs. These techniques include the optimal initial setting of MRs resonant wavelength, increasing the 3-dB bandwidth of optical switching elements by parallel coupling multiple MRs, and the use of passive-routing optical router to minimize the number of switching stages in the mesh-based ONoCs. In this paper, we use a $8 \times 8 \times 2$ 3-D mesh-based ONoC as a case study and analyze the thermal-induced power overhead based on the ONoC thermal model proposed in [12].

III. Optical Routers for 3-D Mesh-Based ONoCs

A. 3-D Mesh-Based ONoC

Fig. 1(a) shows the topology of a 3-D mesh-based ONoC. The proposed 3-D mesh-based ONoC 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 are transmitted along the reserved optical path after the path setup. High-speed optical transmission is achieved without buffering in intermediate routers. *XYZ* dimension-order routing is used for path selection. It is a low-complexity distributed algorithm without any routing table, and is particularly suitable for the mesh and torus-based NoCs. Each processor is assigned a unique ID of (x_i, y_i, z_i) for addressing, and the local router has the same address. For a packet from the source processor (x_s, y_s, z_s) to the destination (x_d, y_d, z_d), it would be first routed in the X dimension until to the router with an address of (x_d, y_s, z_s), and then along the perpendicular Y dimension to the router with an address of (x_d, y_d, z_s). If the destination is in the other layer, that is, $z_s \neq z_d$, the packet would then get to the destination along the Z dimension.

For a packet transmission, the router control unit at the source will generate a path-setup packet to reserve optical links for the payload transmission. The optical switching fabrics of the intermediate routers will be configured properly, where the MR granularity is reached during the reservation. Every time the path-setup packet progresses to the next router, the optical waveguide and electronic interconnect between the previous and current router will also be reserved. A reservation table is used in each router control unit to identify the state of the local router ports. In case the target optical link has been reserved by another transmission, the path-setup packet will be stalled at the current router and wait for the release of the link. Deadlock-free is guaranteed by using the dimension-order routing. We used a new protocol, called quickly acknowledge and simultaneously tear-down (QAST), to reduce control delays during the path setup and tear-down processes. If the path reservation is successful, an optical acknowledgement signal would be generated by the destination and be transmitted back along the reserved optical path. As distinguished from the previous designs, QAST utilizes the symmetric property of optical paths to send back the acknowledgement signal instead of using the electronic control network. This can significantly reduce the setup time of an optical path especially before network saturations. Upon receiving the acknowledgement, the source processor will pass the payload to the optical/electronic (O/E) interface. At the same time, a single-flit tear-down packet will also be sent to the destination through the electronic control network. When the path-setup packet routes in the electronic control network, it not only reserves the optical resources on the path, but also reserves the electronic path in the control network. The tear-down packet is delivered along the successfully prereserved path in the electronic control network. It contains a time-to-live (TTL) field that indicates the necessary number of clock cycles for the payload transmission. With known transmitted packet size L_{packet}, optical interconnect bandwidth B_{optical}, and control network frequency f_{control}, the number of cycles $C_{\text{transmission}}$ required for the optical transmission stored in the TTL field can be easily calculated as (1). The number in the TTL field will be decreased with elapsed cycles by each router along the optical path. Upon receiving the tear-down packet, the router control unit will set the corresponding countdown counter based on the TTL field and start the countdown immediately. Resources associated with the transmission will be released when the countdown counter is timeout. Compared with the traditional tear-down procedure, QAST sends a tear-down packet at the beginning of a transmission instead of at the end of the transmission, which helps reuse the network resources more efficiently. The QAST protocol is important to reduce the packet delays, which is also

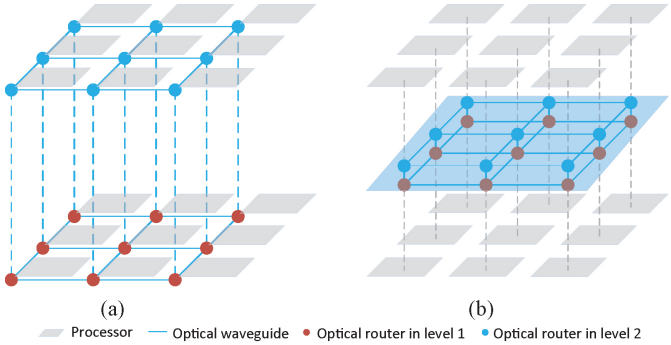

Fig. 1. Topology and floorplan of the 3-D mesh-based ONoC.

a fundamental performance metric besides the throughput

$$C_{\text{transmission}} = \lceil \frac{L_{\text{packet}} \cdot f_{\text{control}}}{B_{\text{optical}}} \rceil. \quad (1)$$

We proposed an adaptive power control mechanism (APC) for the 3-D mesh-based ONoCs to improve the power efficiency. The adaptive power control mechanism is based on the following observations. In ONoCs, power dissipated in an O/E interface is mainly governed by the laser source. For example, in an 80 nm design, while the O/E interfaces consume about 2.5 pJ/bit, the laser sources consume about 1.68 pJ/bit, which accounts for a large proportion of the total O/E power consumption [30]. Traditionally, in order to guarantee enough power for all the possible transmissions, the worst-case optical power loss in an ONoC is considered, and the laser sources are set to offer the worst-case optical power for all the packets. This causes unnecessary power consumption in most transmissions, and also causes the destination circuits to receive optical power within a large dynamic range. The adaptive power control mechanism uses routing information to calculate the optical power loss encountered on an optical path and control the laser source to generate just-enough optical power for the transmission. Regarding laser source, VCSELs are an attractive candidate because of their low power consumption, manufacturing advantages, and high modulation bandwidth. Since they are surface emitting, high-volume production in arrays are possible. VCSELs operating up to 10–40 Gb/s have been reported in [29], [31], and [32]. This makes it possible to achieve high-speed optical transmissions. The light from VCSELs can be coupled to planar waveguides by waveguide grating couplers. The VCSEL diameter is about 55 μm. In an $8 \times 8 \times 2$ mesh-based 3-D ONoC, 128 VCSELs need a total area of 0.4 mm^2. VCSELs provide an opportunity for a better integration and are used by many ONoC architectures to fully integrate ONoCs on chip multiprocessors. The 3-D mesh-based ONoC proposed in this paper can also work with an off-chip laser.

In the 3-D mesh-based ONoCs with dimension-order routing, an optical path is only determined by the source and destination addresses, and the optical power loss on different optical paths can be calculated. A precalculated table can be used. While the router control unit tries to setup an optical path, the adaptive power control unit will calculate the minimum laser power required and control the VCSEL driver. The minimum laser power needed can be calculated by adding the optical power loss of the path to the minimum optical power required at the receiver. We assume to use the 3-D integration technology to connect VCSELs with the underlying CMOS driver circuits through TSVs. The output power of VCSELs is directly modulated by the driving current, instead of doing the modulation by distributed modulators. The adaptive power control mechanism can be implemented by changing the voltage level of the VCSEL drivers. Compared to nonadaptive mechanisms, the adaptive power control mechanism saves the dynamic power consumption of VCSELs and improves the ONoCs power efficiency. The transmitter power-on delay includes the VCSEL driver circuit delay and the VCSEL device turn-on delay. We have considered the transmitter power-on delay in the following simulations. It is in parallel with the procedure of the path setup.

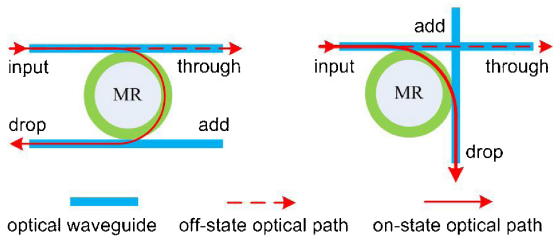

Fig. 2. Basic 1×2 switching elements.

B. New Optical Routers

Based on the switching functions of MRs, several 5×5 optical routers have been proposed for routings in the mesh or torus-based ONoCs [22]–[25]. For the 3-D mesh-based ONoC we proposed, 5×5 optical routers are not enough for the connection requirements. Optical routers with 6×6 and 7×7 switching functions are required in the 3-D mesh networks. To the best of our knowledge, this is the first paper that proposes compact passive-routing 6×6 and 7×7 optical routers for dimension-order routing in the 3-D mesh-based ONoCs. We also show the traditional crossbar-based 7×7 optical routers in this section. Comparisons show that the new proposed optical routers cost much less than the crossbar-based optical routers. Besides, we proposed to use 4×4 and 5×5 optical routers at the network edges which further minimizes the cost of optical resources.

The optical switching fabric is built from two types of 1×2 basic switching elements (Fig. 2), including the crossing element and the parallel element, both of which consist of one MR and two optical waveguides. The crossing element has a waveguide crossing that will introduce a crossing insertion loss for passing signals. Although the loss per crossing is small, a large number of crossings in the optical transmission path may lead to a significant power loss. The parallel element is preferred in router designs in order to minimize the number of waveguide crossings encountered by the optical signals. The basic elements achieve 1×2 optical switching functions by powering on/off the MR. When powered on, the MR has an on-state resonance wavelength of λ_{on}. For an optical signal that is injected to the input port with a center wavelength of λ_{on}, if the MR is powered on, it will be coupled into the MR and be delivered to the drop port; otherwise, the optical signal will propagate directly to the through port.

Figs. 3 and 4 show the 6×6 and 7×7 strictly nonblocking low-loss optical routers for dimension-order routing in the 3-D mesh-based ONoCs. The router architectures are designed to minimize the number of MRs and switching activities. With certain device technologies, the power efficiency of a ONoC is mainly governed by the optical power loss that is encountered by the light signal along the path. Waveguide crossings in ONoCs do not affect the bandwidth, but cause more optical power loss and power consumption in packet transmissions. Each waveguide crossing introduces about 0.12 dB insertion loss to the passing optical signals [21]. In order to minimize

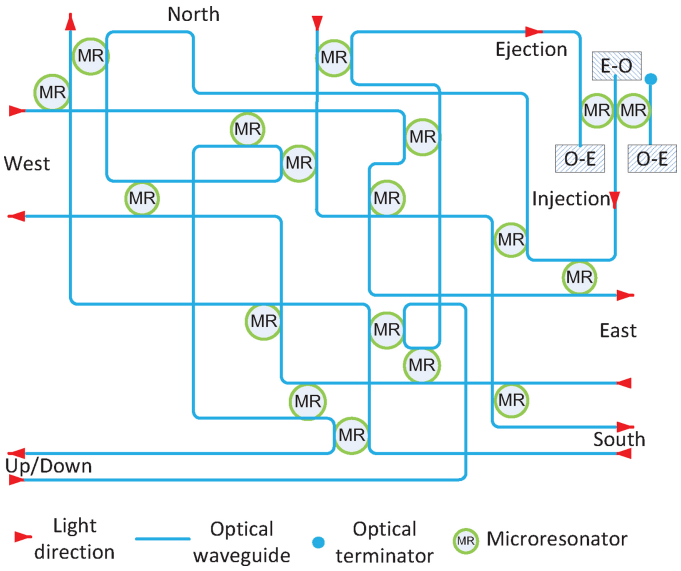

Fig. 3. 6 × 6 optical router for dimension-order routing in 3-D mesh-based ONoC.

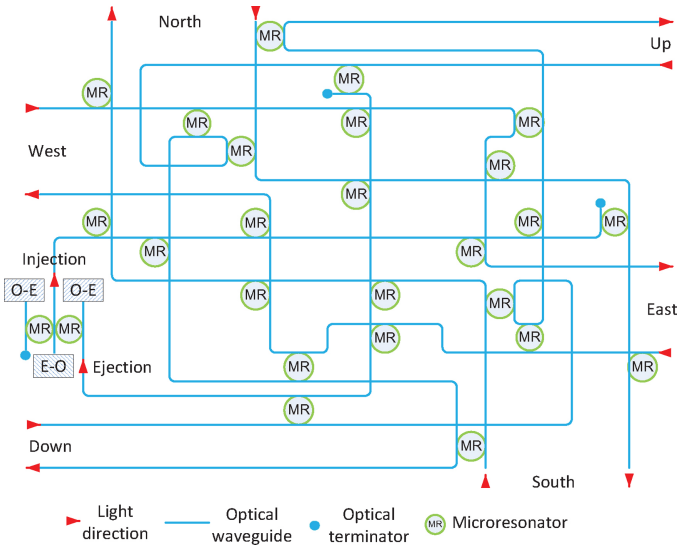

Fig. 4. 7 × 7 optical router for dimension-order routing in 3-D mesh-based ONoC.

power loss in ONoCs, the router architectures are designed to minimize the number of waveguide crossings.

The 6 × 6 router uses only 18 MRs and one optical terminator. It has six bidirectional ports, namely East, West, South, North, Up (or Down) and Injection/Ejection. The ports are aligned to their intended directions, and the input and output of each port is also properly aligned to ensure that no extra waveguide crossing is caused while one 6 × 6 router is connected to another in the 3-D mesh network. The nonblocking property of this router is proved by enumerating all possible cases. The number of waveguide crossings is 19 in the 6 × 6 router. The internal structure of the 6 × 6 router is designed to passively route packets. Packets traveling between the East and West, South and North, do not require powering on any MR. Besides, only one MR is required to be powered on when the packet makes a turn. The 7 × 7 router is designed

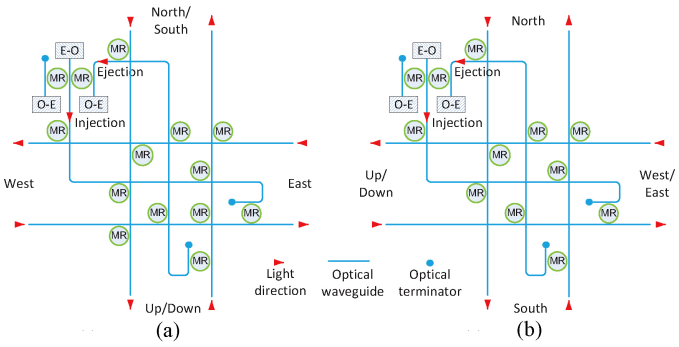

Fig. 5. 5 × 5 optical router Crux for dimension-order routing in 3-D mesh-based ONoC. (a) 5 × 5 optical router Crux. (b) Modified Crux.

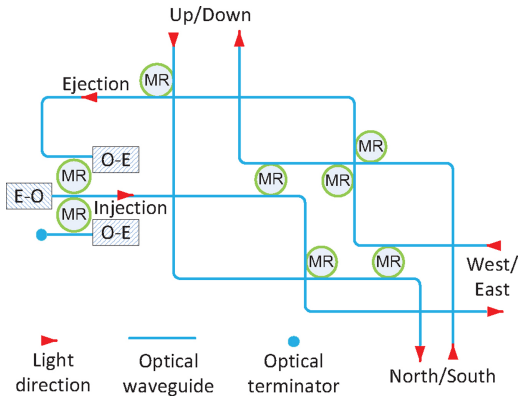

Fig. 6. 4 × 4 optical router for dimension-order routing in 3-D mesh-based ONoC.

based on the 6 × 6 router architecture. The 7 × 7 router uses 26 MRs and three optical terminators. It inherits the passive routing property of the 6 × 6 router proposed above. This feature is beneficial to the scalability of a network. It also guarantees that the maximum power consumption to route a packet through a dimension-order routing network is a small constant number.

In the 3-D mesh-based ONoCs, routers at the network edges or corners do not fully utilize the 6 × 6 or 7 × 7 optical switching functions. For routers at the network edges, a 5 × 5 optical switching function is enough to satisfy the connection requirements. In order to reduce network resources, passive-routing 5 × 5 optical routers Crux and the modified Crux (Fig. 5) [23], are proposed to be used at the network edges. Crux uses 14 MRs and the modified Crux uses two less MRs. Besides, we designed a 4 × 4 optical router (Fig. 6) for the network corners of the 3-D mesh-based ONoC. The 4 × 4 optical router only uses eight MRs and has eight waveguide crossings inside.

To implement the optical acknowledgement mechanism, an additional O/E interface is needed in case the O/E interfaces of the communication pair are not available to send or receive the optical acknowledgement. In the reserved optical path, only one direction is used at any given time. With the additional O–E receiver and MRs, no interference would be caused by the optical acknowledgement transmission. As shown in Figs. 3 and 4, a small circuit of O/E interface is used in every optical router. It uses one additional O–E receiver and two additional

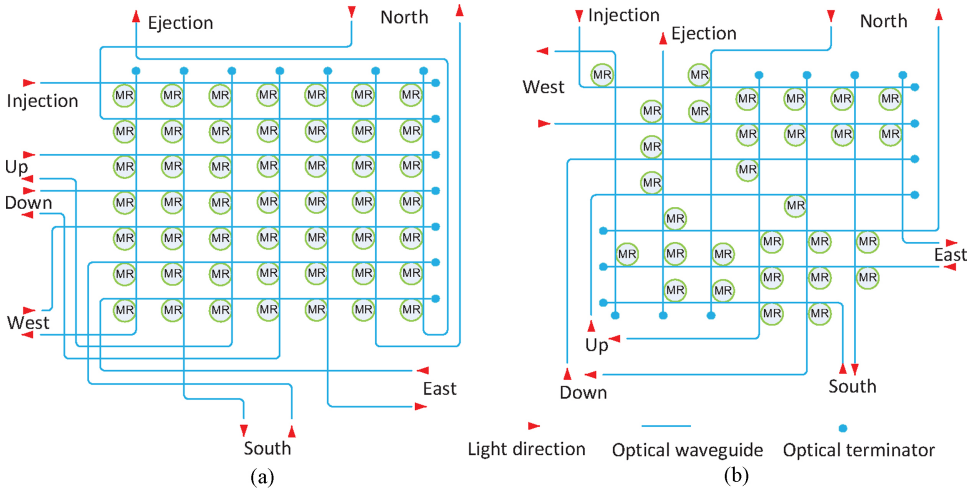

Fig. 7. Crossbar based 7 × 7 optical routers. (a) Traditional crossbar based 7 × 7 optical router. (b) Optimized partial crossbar based 7 × 7 optical router.

TABLE I
NUMBER OF MRS AND WAVEGUIDE CROSSINGS IN DIFFERENT ROUTERS

Optical Router	N_{MR}	$N_{crossing}$	$N_{terminator}$
4 × 4	8	8	1
5 × 5	14 (or 12)	11	3
6 × 6	18	19	1
7 × 7	26	31	3
Traditional 7 × 7 crossbar	49	67	14
Partial 7 × 7 crossbar	30	47	14

MRs, without extra E–O transmitter. One MR is used to couple an optical signal from the E–O transmitter into the ejection port. In this way, even if the E–O transmitter is occupied by sending the payload, an optical acknowledgement signal can still be sent out by coupling a portion of power to the ejection port without interrupting the correctness of the data. The portion of power transferred to the drop port can be designed before fabrication by properly choosing device parameters (e.g., the waveguide power coupling coefficients, bending). The adaptive power control mechanism is used to adjust the optical power in the E–O transmitter, in case it is going to send an optical acknowledgement while sending the payload data. The E–O transmitter has two working modes. It works at the normal mode at most of times, and works at the double-power mode if it sends the optical acknowledgement and payload data at the same time. In the double-power mode, the adaptive power control mechanism adjusts the optical power in the transmitter to offer the additional power consumed by the optical acknowledgement. Another MR is used for the source router to sense the optical acknowledgement signal. After the source sends out a path-setup request, it will turn on this MR and then listen to the signals coming from the injection port. When the optical acknowledgement arrives, it would be coupled by the MR to the additional O–E receiver, without disturbing the original O–E receiver which might be receiving payload data.

Optical routers from 4 × 4 to 7 × 7 are designed for dimension-order routing in the 3-D mesh-based ONoCs. Table I shows the number of MRs and the total number of waveguide crossings in different routers. We compared the 7 × 7 router with the crossbar based 7 × 7 optical routers (Fig. 7). The optimized partial crossbar is reduced from the traditional crossbar for dimension-order routing in 3-D meshes. Different from the passive routing feature of the new proposed optical routers, one MR is required to be turned on for every switching in the crossbar-based optical routers. Even with the additional MRs for the optical acknowledgement mechanism, the new proposed 7 × 7 router uses 23 less MRs than the traditional crossbar and four less MRs than the optimized partial crossbar. The new 7 × 7 router has much less optical waveguide crossings than the crossbar-based routers. It is expected that the new 7 × 7 router will have a better power loss than the crossbar-based optical routers.

IV. FLOORPLAN OF THE 3-D MESH-BASED ONOCS

A. Floorplan With Routers in a Single Optical Layer

In order to minimize the number of optical layers, we propose to merge two or more optical layers into a single optical layer [Fig. 1(b)] while processor cores are still implemented in two electrical layers. We assume that the two electrical layers are stacked together with TSVs. The optical layer is stacked on the top of the two electrical layers. Optical devices can be accessed by the two electrical layers through TSVs.

When merging the 3-D optical routing network to one single optical layer, optical waveguides from different levels would intersect with each other and extra waveguide crossings would be caused. In order to minimize the optical power loss in ONoCs, the floorplan of the optical routing layer should be designed to reduce the total number of waveguide crossings in the floorplan. Fig. 8(a) shows a floorplan design for the 3-D $M \times N \times 2$ mesh-based ONoCs. It follows the regular 3-D mesh topology but implements all optical routers in one optical layer. We find that this floorplan design works well for reducing the number of waveguide crossings, especially for large-sized 3-D $M \times N \times 2$ ($M, N \geq 8$) mesh-based ONoCs. Table II shows the number of waveguide crossings in the floorplan of different sizes, excluding the waveguide crossings inside the optical

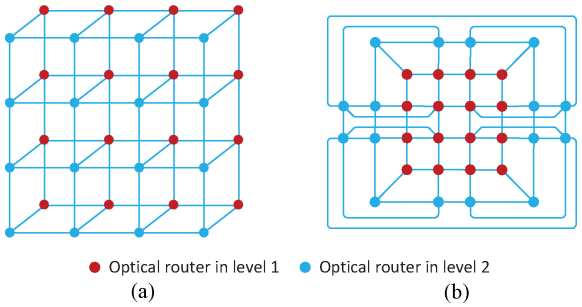

Fig. 8. Floorplan of 3-D M×N×2 mesh-based ONoC, with routers in a single optical layer. (a) Floorplan with routers in a single optical layer. (b) Floorplan after special optimization.

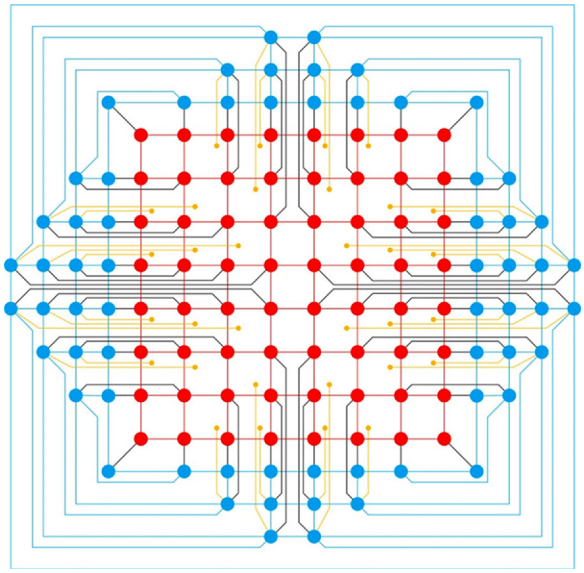

Fig. 9. Optimized floorplan of 8 × 8 × 2 mesh-based ONoC.

routers. We assumed the optical links are all bidirectional and there are two optical waveguides for each link in the floorplan.

B. Special Floorplan Optimization

For small-sized $M \times N \times 2$ 3-D mesh-based ONoCs ($M, N \leq 8$), the floorplan in Fig. 8(a) may not be the optimal physical floorplan. In order to reduce waveguide crossings in the network level, we optimized the floorplan for the small-sized $M \times N \times 2$ 3-D mesh-based ONoCs ($M, N \leq 8$) by rearranging the optical waveguides. The optimized floorplan maintains the connection property shown in the $M \times N \times 2$ 3-D mesh topology, but minimizes the number of waveguide crossings in the physical implementation. Fig. 8(b) shows the floorplan of a 4 × 4×2 3-D mesh-based ONoC after the optimization. After the floorplan optimization, the number of waveguide crossings in the floorplan is reduced from 36 [Fig. 8(a)] to 16.

Following the rules, the floorplan for an 8 × 8 × 2 3-D mesh-based ONoC is designed as shown in Fig. 9. Though we merged the two routing layers into one optical layer, we assumed that the processors are still placed in two electrical layers. For routers in the outlying area of the optical layer, corresponding processors are located in appropriate locations

TABLE II
NUMBER OF WAVEGUIDE CROSSINGS IN THE FLOORPLAN OF AN M×N×2 MESH-BASED ONOC

M×N×2	Floorplan in Fig. 8(a)	Special Floorplan Optimization
4 × 4×2	36	16
5 × 5×2	64	40
6 × 6×2	100	80
7 × 7×2	144	140
8 × 8 × 2	196	224

TABLE III
PARAMETERS ASSUMED FOR SIMULATION

Optical link bandwidth	40 Gb/s
Electronic clock frequency	1.25 GHz
Metallic interconnect width	32-b bidirectional
CMOS technology	45 nm

to minimize the area of the second processor layer. In Fig. 9, yellow dots represent the location of processors which are connected to routers in outlying loops.

Table II shows the comparison of the number of waveguide crossings in two kinds of floorplans, excluding the waveguide crossings inside the optical routers. We can conclude that when $M, N < 8$, the special floorplan optimization has less waveguide crossings than the floorplan design in Fig. 8(a); when $M, N \geq 8$, the floorplan design in Fig. 8(a) would work better.

V. SIMULATION RESULTS AND COMPARISON

The motivation of this paper is to enhance the performance and energy efficiency by taking advantage of 3-D mesh over 2-D mesh, together with the new low-cost optical routers and floorplan optimizations. So we use the 2-D mesh-based NoCs as a baseline for comparison. All the comparisons between different NoCs are on equal footing for the same number of processor cores. We evaluated the performance and energy efficiency of the proposed 3-D mesh-based ONoC for MPSoCs in 45 nm, and compared it with the matched 2-D mesh-based ONoC and the 2-D mesh-based electronic NoC. Using the optical network for all control packets may further improve the performance for the 3-D mesh-based ONoC. We may consider this in future works. SystemC-based cycle-accurate simulators are developed for network simulations of the proposed 3-D mesh-based ONoC and the referenced NoCs. We assumed the same 40 Gb/s data-link bandwidth for all the comparing NoCs in this paper. For the proposed 3-D mesh-based ONoC, the electronic control network operates at 1.25 GHz with 32-b-wide bidirectional metallic interconnects (Table III). The electronic control network works according to the protocols described in Section III-A. All the control delays such as path setup delay are included in network simulations. The electronic control unit works at 1.25 GHz clock frequency and the power consumption is about 12.625 mW. The total silicon area of the 8 × 8 × 2 mesh-based electronic control network is about 2.15 mm^2, including the router control units, laser drivers, and receiver circuits. The matched 2-D mesh-based ONoC also uses circuit switching with an electronic

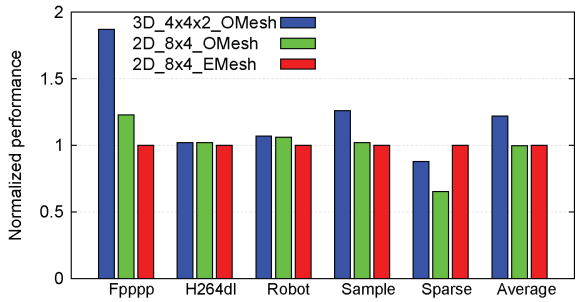

Fig. 10. Normalized performance comparison under real applications, 32-core.

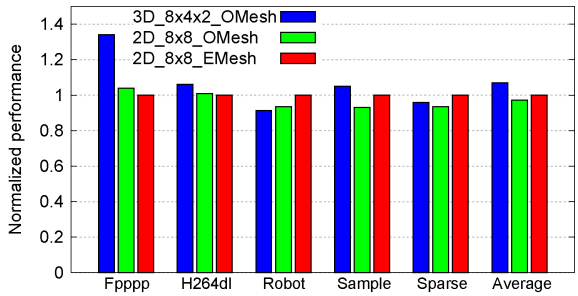

Fig. 11. Normalized performance comparison under real applications, 64-core.

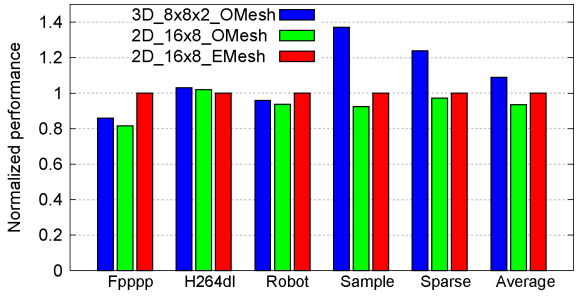

Fig. 12. Normalized performance comparison under real applications, 128-core.

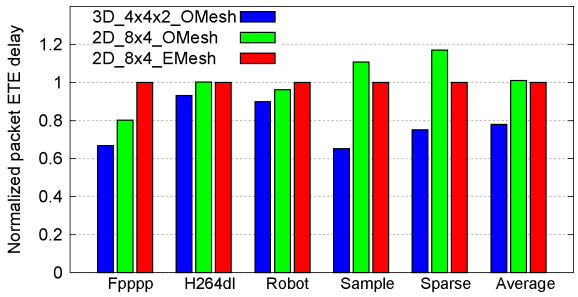

Fig. 13. Normalized average packet ETE delay under real applications, 32-core.

control network. The electronic control network is also designed to work at 1.25 GHz clock frequency with 32-b-wide bidirectional interconnects. It employs 5×5 optimized partial crossbars as the optical switching fabric. For the matched 2-D mesh-based electronic NoC, electronic worm-hole switching is adopted and two virtual channels (each virtual channel is 16-flit deep) are used in each input port. The 2-D electronic mesh-based NoC also works at 1.25 GHz, and 40 Gb/s bandwidth for each port is offered by using 32-b-wide bidirectional metallic interconnects.

The simulations are based on a set of real applications, including Fpppp, H264 decoder, Sample, Robot, and Sparse [33]. For a massive processing of data streams, a large-scale multiprocessor system would be required for performance and power efficiency requirements. Meanwhile, an efficient communication architecture is needed to guarantee that data are delivered in time. As distinguished from random traffic models, real applications have fixed access patterns and tend to have nearest neighboring communication patterns [34]. Before the network simulations, an offline optimization approach was applied for each application to map and schedule tasks onto the MPSoC with the objective of maximizing the system performance [35]. Besides the real applications, we also evaluated the network performance of the 3-D mesh-based ONoC under uniform traffic. For the uniform traffic pattern, functional cores are assumed to generate packets independently and the packet generation intervals follow a negative exponential distribution.

A. Performance Comparison

1) *Real Applications:* We simulated the 3-D mesh-based ONoC for 32-core, 64-core, and 128-core MPSoCs under the real applications. Figs. 10–12 show the normalized performance comparisons under different applications. For the 32-core MPSoC, we compared the performance of the proposed $4 \times 4 \times 2$ 3-D mesh-based ONoC with the 8×4 2-D mesh-based ONoC, as well as with the 8×4 2-D mesh-based electronic NoC. On average of the set of real applications, the 3-D $4 \times 4 \times 2$ mesh-based ONoC achieves about 22% and 23% performance improvements, respectively, compared with the 2-D mesh-based electronic NoC and ONoC. For the 64-core MPSoC, the proposed 3-D ONoC is in an $8 \times 4 \times 2$ 3-D mesh, and the 2-D NoCs are in an 8×8 2-D mesh. On average of the set of real applications, the 3-D $8 \times 4 \times 2$ mesh-based ONoC achieves about 7% and 10% performance improvements, respectively, compared with the 2-D mesh-based electronic NoC and ONoC. For the 128-core MPSoC, the proposed 3-D ONoC is based on an $8 \times 8 \times 2$ 3-D mesh, and the 2-D NoCs are in a 16×8 mesh. The 3-D $8 \times 8 \times 2$ mesh-based ONoC achieves about 9% and 17% performance improvements, respectively, compared with the 2-D mesh-based electronic NoC and ONoC.

Figs. 13–15 show the normalized average packet end-to-end (ETE) delay comparisons under different applications. It is shown that the proposed 3-D mesh-based ONoC improves the average packet ETE delay over the 2-D counterparts. For the 32-core MPSoC, on average of the set of applications, the 3-D $4 \times 4 \times 2$ mesh-based ONoC reduces 22% and 23% of average packet ETE delay compared with the 2-D mesh-based electronic NoC and ONoC. For the 64-core MPSoC, on average of the set of applications, the 3-D $8 \times 4 \times 2$ mesh-based ONoC reduces 8% and 10% of average packet ETE delay compared with the 2-D mesh-based electronic NoC and ONoC. For the 128-core MPSoC, on average of the set of applications, the 3-D $8 \times 8 \times 2$ mesh-based ONoC reduces 7% and 9% of average packet ETE delay compared with the 2-D mesh-based electronic NoC and ONoC.

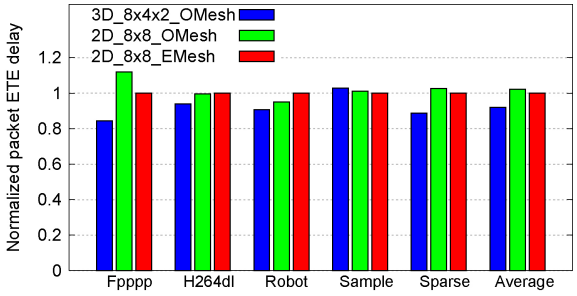

Fig. 14. Normalized average packet ETE delay under real applications, 64-core.

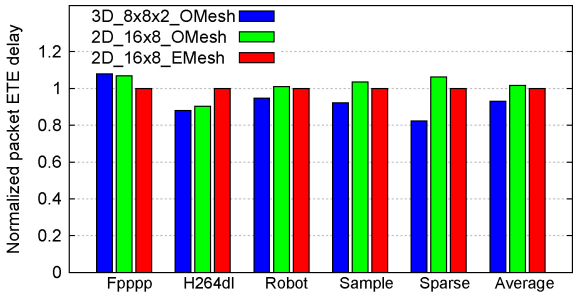

Fig. 15. Normalized average packet ETE delay under real applications, 128-core.

2) *Uniform Traffic:* We simulated the 3-D mesh-based ONoC for 32-core, 64-core, and 128-core MPSoCs under the uniform traffic. Figs. 16–18 show the simulation results with packet size of 512 bytes. The network performance was evaluated and compared in metrics of the average packet ETE delay and the network throughput. The packet ETE delay is the average time a packet takes to reach the destination, and the network throughput is defined as the total data transfer rate in the network under a given injection rate. The injection rate is defined as ((2)), where T_{transmit} is the time to transmit the packet and T_{interval} is the average time interval between the generations of two successive packets. T_{interval} follows a negative exponential distribution. For example, if the injection rate is 0.5, the average time interval between the generations of two successive packets is equal to the time of one packet transmission

$$\alpha = \frac{T_{\text{transmit}}}{T_{\text{transmit}} + T_{\text{interval}}}. \tag{2}$$

As shown in Figs. 16–18, the 3-D mesh-based ONoC outperforms the 2-D mesh-based NoCs for both the network throughput and the average packet delay. The network throughput increases with the injection rate, and after a saturation point, the throughput stops increasing. The 3-D mesh-based ONoC achieves a better throughput than the 2-D mesh-based NoCs. With 512 B packets, the 4 × 4×2 3-D mesh-based ONoC has a saturation throughput of 530 Gb/s. The 2-D 8×4 mesh-based ONoC and electronic NoC saturate at an earlier injection rate with a saturation throughput of about 359 Gb/s and 394 Gb/s, respectively. For the 64-core MPSoC, the 8 × 4 × 2 3-D mesh-based ONoC saturates at a throughput of 694 Gb/s. It is higher than the 8×8 2-D optical mesh and electronic mesh by about 138 Gb/s and 50 Gb/s, respectively. For the 128-core MPSoC, the saturation throughput of the

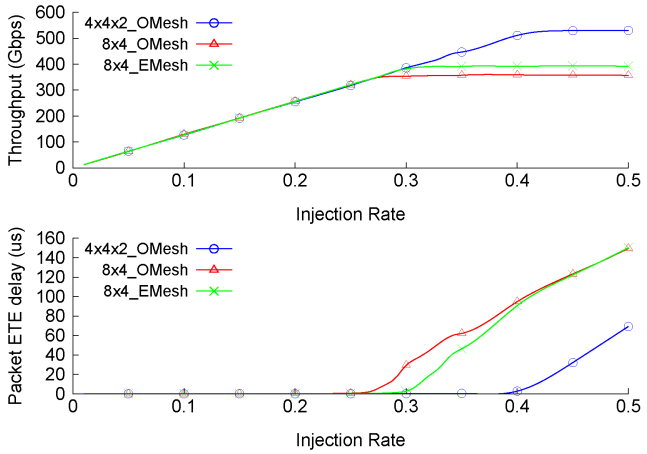

Fig. 16. Performance comparison of different 32-core ONoCs under uniform traffic, with 512-byte packets.

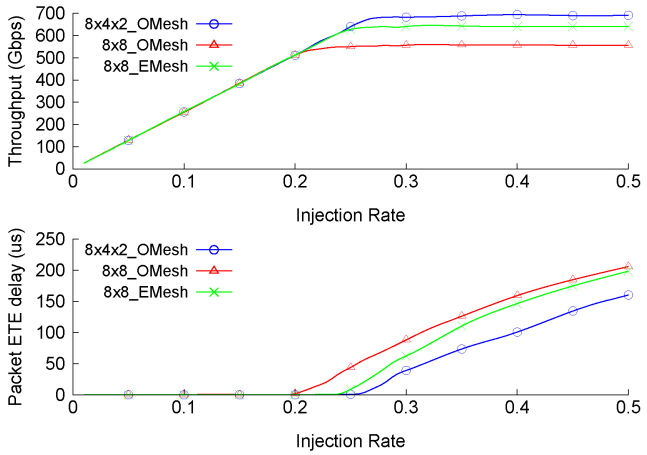

Fig. 17. Performance comparison of different 64-core ONoCs under uniform traffic, with 512-B packets.

8 × 8 × 2 3-D mesh-based ONoC increases to 1069 Gb/s. It is much higher than the 16×8 2-D mesh-based ONoC and electronic NoC by 444 and 313 Gb/s, respectively.

As shown in Figs. 16–18, the 3-D mesh-based ONoC has the best packet delays. For example, when the injection rate is 0.03, the average packet delay of the 4 × 4×2 3-D mesh-based ONoC is about 0.114 μs, while the 8 × 4 2-D mesh-based ONoC and electronic NoC have an average packet delay of 0.118 μs and 0.126 μs, respectively. After the network saturations, the packet delay of all the three NoCs increases dramatically. The average packet delay of the 4 × 4×2 3-D mesh-based ONoC increases to 32.189 μs at the injection rate of 0.45. At the same injection rate, the 8 × 4 2-D mesh-based ONoC and electronic NoC have a much larger packet delay of 123.117 μs and 122.267 μs, respectively.

B. Energy Efficiency Comparison

Power consumption is a critical aspect in NoC designs. For high-performance computing, low power consumption can reduce the cost related to packaging, cooling solution, and system integration. With technology scaling, the on-chip communication demands an increasing proportion of the system

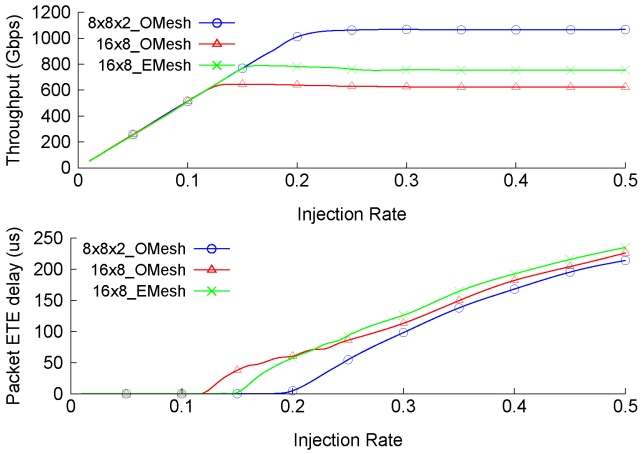

Fig. 18. Performance comparison of different 128-core ONoCs under uniform traffic, with 512-B packets.

power budget. In current prototypes with tens of cores, the power consumed by the electronic NoC accounts for over 25% of the overall power, which is too high to meet the expected requirements of future MPSoC systems [36]. The introduction of optical interconnects helps manage the power budget in multicore processor architectures [37], but a better optical power control mechanism is still desired for further energy saving. We evaluated the energy efficiency of the 3-D mesh-based ONoC with the adaptive power control (APC), and compared it with the 2-D mesh-based ONoC as well as the electronic one under the set of real applications and uniform traffic pattern. We also assume that the 2-D mesh-based ONoC is equipped with the adaptive power control mechanism.

The 3-D mesh-based ONoC is a circuit-switched network with an electronic control network for path maintenances. The power consumption for a packet transmission involves the energy consumed for payload transmission in the optical domain and the energy consumed for routing control packets in the electronic domain. The energy consumption for a control packet is estimated as the energy required to transfer it through all the electrical interconnects, electronic switching fabrics, and control units along the path in the electronic control network. The energy consumed by a payload transmission includes the energy consumed by the O/E interfaces, and MRs energy consumption in the optical path. Since the size of the control packets is small, the power consumed in the electronic control network only takes a small proportion, and the power consumption in O/E interfaces accounts for the major proportion of the total energy consumption.

A typical O/E interface includes a serializer/deserializer, laser driver, VCSEL, optical waveguide, photodetector, and TIA-LA circuits. For the evaluation of O/E conversions power consumption, we use the VCSEL model in [29], the serializer and deserializer designs in [38], and the VCSEL driver and TIA-LA circuit designs in [30]. Since the electronic part of the 3-D mesh-based ONoC is in 45 nm, we scale all the related power consumption linearly to 45 nm. For example, the driver and TIA-LA circuits power consumption is scaled from 0.2 and 0.6 pJ/bit in 80 nm to 0.1125 and 0.3375 pJ/bit in 45 nm,

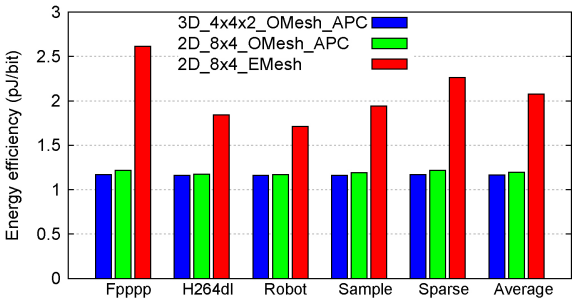

Fig. 19. Average energy efficiency of 3-D $4\times4\times2$ mesh-based ONoC under real applications.

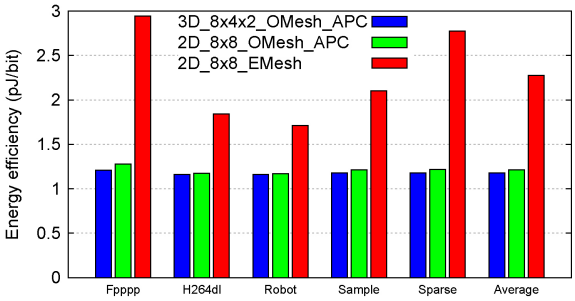

Fig. 20. Average energy efficiency of 3-D $8\times4\times2$ mesh-based ONoC under real applications.

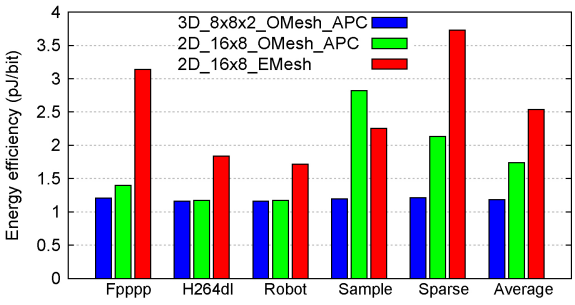

Fig. 21. Average energy efficiency of 3-D $8\times8\times2$ mesh-based ONoC under real applications.

and the power consumption of the serializer and deserializer is scaled from 0.576 pJ/bit in 90 nm to 0.288 pJ/bit in 45 nm. The photodetector model is based on a 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) [39]. Besides, the power consumption for turning on a MR is assumed to be 20 μW.

Figs. 19–21 compare the energy efficiency of the 3-D mesh-based ONoC with the matched 2-D mesh-based ONoC and electronic NoC for 32-core, 64-core, and 128-core MPSoCs in 45 nm. It is shown that the 3-D mesh-based ONoC achieves the best energy efficiency for all the applications. For the 32-core MPSoC, on average of the set of applications, the 3-D $4\times4\times2$ mesh-based ONoC reduces 44% of energy consumption compared with the 2-D mesh-based electronic NoC, and has a slight improvement over the 2-D mesh-based ONoC. For the 64-core MPSoC, the 3-D $8\times4\times2$ mesh-based ONoC reduces 52% of energy consumption compared with the 2-D mesh-based electronic NoC. For the 128-core MPSoC, the improvement becomes more obvious. The 3-D

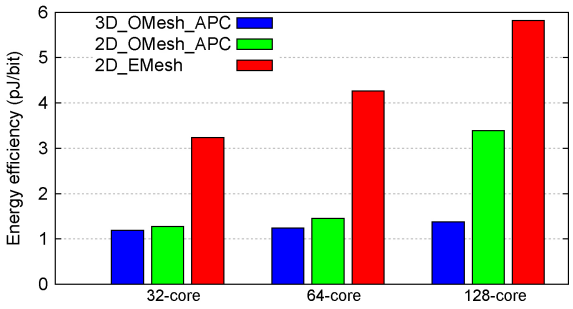

Fig. 22. Average energy efficiency of NoCs under uniform traffic, with 512-B packets.

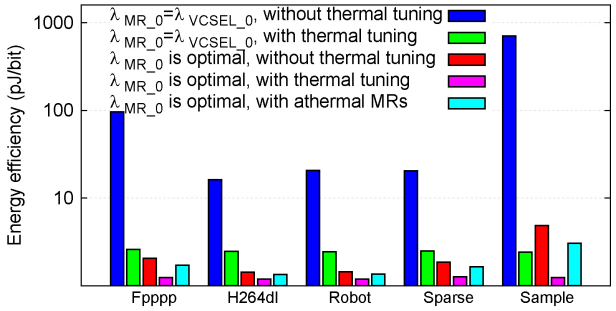

Fig. 23. Energy efficiency of 3-D $8\times8\times2$ mesh-based ONoC under different real applications, assuming the maximum chip temperature reaches 85 °C.

$8 \times 8 \times 2$ mesh-based ONoC reduces 47% and 32% of energy consumption compared with the 2-D mesh-based electronic NoC and the 2-D mesh-based ONoC.

Fig. 22 shows the comparison of energy efficiency under the uniform traffic. It shows that for the 32-core MPSoC, the 3-D $4 \times 4 \times 2$ mesh-based ONoC reduces 63% of energy consumption compared with the 2-D mesh-based electronic NoC, and has a slight improvement over the 2-D mesh-based ONoC. For the 64-core MPSoC, the 3-D $8 \times 4 \times 2$ mesh-based ONoC reduces 71% and 14% of energy consumption compared with the 2-D mesh-based electronic NoC and the 2-D mesh-based ONoC. For the 128-core MPSoC, the 3-D $8 \times 8 \times 2$ mesh-based ONoC reduces 76% and 59% of energy consumption compared with the 2-D mesh-based electronic NoC and the 2-D mesh-based ONoC.

$$10\log((I-\alpha-\beta(T_{\mathrm{VCSEL}}-T_{th})^2)(\varepsilon-\gamma\cdot T_{\mathrm{VCSEL}})) - \sum_{i=1}^{N} 10\log((\frac{2\kappa^2+\kappa_p^2}{2\kappa^2})^2\cdot(1+\delta^{-2}(\lambda_{\mathrm{VCSEL\_min}} + \rho_{\mathrm{VCSEL}}(T_{\mathrm{VCSEL}}-T_{\min}) - \rho_{\mathrm{MR}}(T_{\mathrm{MR}_i}-T_{\min}) - \lambda_{\mathrm{MR\_}min})^2)) - L_{\mathrm{WG}} \geq S_{\mathrm{RX}}. \quad (3)$$

Thermal sensitivity is a potential issue in ONoC designs. As a result of thermo-optic effect, the temperature-dependent wavelength shifts in VCSELs and silicon-based MRs are found to be about 50–100 pm/°C [27], [28]. As a widely used device in ONoCs, MRs perform as a wavelength-selective optical switch or modulator. The thermal related wavelength variations will result in additional optical power loss. Besides, the VCSEL power efficiency degrades at high temperatures [29]. Based on the system-level analytical ONoC thermal model proposed in [12], we quantitatively study the thermal effects in the 3-D $8 \times 8 \times 2$ mesh-based ONoC based on the low-cost optical routers. The passive routing feature of the new proposed optical routers guarantees that the maximum number of switching stages in the 3-D $8 \times 8 \times 2$ mesh-based ONoC with dimension-order routing is four. To ensure that ONoCs function properly, a necessary condition is that the optical signal power received by a receiver should not be lower than the receiver sensitivity. This condition must hold, otherwise the BER would increase significantly. We assume the photodetector sensitivity is -14.2 dBm for a BER of 10^{-12} [39]. We model the condition in (3), assuming that the temperature range is between $T_{\min}$ and $T_{\max}$ ($T_{\min} \leq T_{\mathrm{VCSEL}}, T_{\mathrm{MR}_i} \leq T_{\max}$) [12]. The temperature-

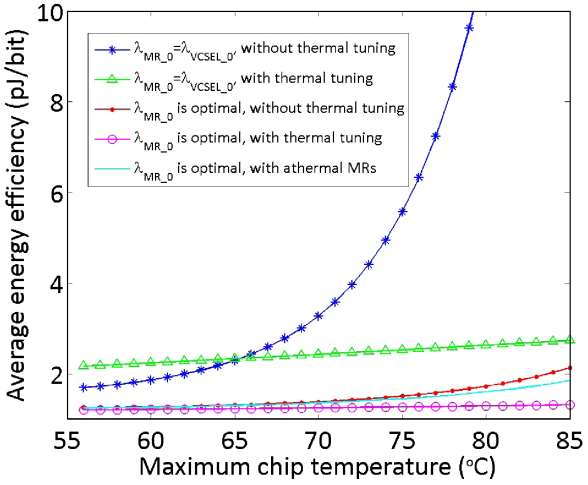

Fig. 24. Average energy efficiency of 3-D $8\times8\times2$ mesh-based ONoC with different configurations.

dependent wavelength shift of VCSELs and MRs are denoted by ρ_{VCSEL} and ρ_{MR}. The first term of (3) is the output power of the VCSEL driven by current I which is above the threshold I_{th}. The second term is the summation of optical power loss due to N switching stages of MRs, assuming that the 3-dB bandwidth of MRs is 2δ. L_{WG} is the waveguide propagation loss in the link, and S_{RX} is the sensitivity of the receiver.

Two traditional techniques have been proposed to compensate the temperature-dependent wavelength shift in MRs, including the active thermal tuning with local microheaters and the passively temperature-compensated athermal MRs [40], [41]. The athermal MRs usually have a limited range of working temperatures, beyond which the MRs would be thermal sensitive. In our case study, we assume that the athermal MRs are working within their required temperature range. Besides the use of low-cost optical routers with a minimized number of switching stages, we proposed a new passive temperature compensation technique using the optimal device setting to improve the power efficiency of ONoCs under thermal variations [12].

We quantitatively evaluate the power efficiency of the 3-D $8 \times 8 \times 2$ mesh-based ONoC under different combinations of device settings and low-temperature-dependence techniques. 55 °C–85 °C (30 °C of variation) is a typical on-chip temperature range, so we use this assumption in our case study. Fig. 23 shows the power efficiency of the 3-D $8\times8\times2$ mesh-

based ONoC under different real applications, assuming that the maximum chip temperature $T_{\max}$ reaches 85 °C. Fig. 24 shows the average power efficiency of the 3-D $8 \times 8 \times 2$ mesh-based ONoC under different maximum temperature $T_{\max}$. We assume the 3-dB bandwidth of MRs is 3.1 nm. If the initial MR resonance wavelength λ_{MR_0} is equal to λ_{VCSEL_0}, the average communication power efficiency is 25 pJ/bit when the maximum chip temperature reaches 85 °C. It can be improved to about 2.7 pJ/bit by applying the traditional thermal tuning technique with microheater, assuming that the tuning efficiency is 3.5 mW/nm [40]. The new technique using the optimal device setting can improve the average power efficiency to 2.1 pJ/bit [12]. It is shown that in this particular case, the effectiveness of the two techniques is comparable. If the two techniques are applied at the same time, the average power efficiency can be further improved to 1.3 pJ/bit. If we apply the optimal device setting with athermal MRs, the average power efficiency is about 1.8 pJ/bit. Since thermal tuning and athermal MRs could impose high fabrication cost, the new technique using the optimal device setting without thermal tuning is also an alternative solution in this particular case.

VI. Conclusion

We proposed a 3-D mesh-based ONoC for MPSoCs, together with the optimized floorplans and low-cost nonblocking 4×4, 5×5, 6×6, and 7×7 optical routers. The 3-D $8 \times 8 \times 2$ mesh-based ONoC achieved 17% performance improvement with 32% less energy consumption than the matched 2-D 16×8 mesh-based ONoC. It achieved 9% performance improvement while reducing 47% of energy consumption as compared to the 2-D 16×8 mesh-based electronic NoC. Additionally, we quantitatively analyzed the thermal effects in the 3-D $8 \times 8 \times 2$ mesh-based ONoC.

References

[1] W. Dally and B. Towles, "Route packets, not wires: On-chip interconnection networks," in *Proc. Des. Autom. Conf.*, 2001, pp. 684–689.

[2] L. Benini and G. De Micheli, "Networks on chip: A new paradigm for systems on chip design," in *Proc. Des. Autom. Test Eur. Conf. Exhibit.*, 2002, pp. 418–419.

[3] N. Kirman, M. Kirman, R. Dokania, J. Martinez, A. Apsel, M. Watkins, and D. Albonesi, "On-chip optical technology in future bus-based multicore designs," *IEEE Micro*, vol. 27, no. 1, pp. 56–66, Jan.–Feb. 2007.

[4] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, "Firefly: Illuminating future network-on-chip with nanophotonics," in *Proc. 36th ISCA*, 2009, pp. 429–440.

[5] K. H. Mo, Y. Ye, X. Wu, W. Zhang, W. Liu, and J. Xu, "A hierarchical hybrid optical-electronic network-on-chip," in *Proc. IEEE Comput. Soc. Annu. Symp. VLSI*, Jul. 2010, pp. 327–332.

[6] A. Shacham, K. Bergman, and L. Carloni, "On the design of a photonic network-on-chip," in *Proc. 1st Int. Symp. Networks-on-Chip*, 2007, pp. 53–64.

[7] 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, and K. Asanovic, "Building manycore processor-to-DRAM networks with monolithic silicon photonics," in *Proc. 16th IEEE Symp. High Performance Interconnects*, Aug. 2008, pp. 21–30.

[8] A. W. Topol, D. C. L. Tulipe, L. Shi, D. J. Frank, K. Bernstein, S. E. Steen, A. Kumar, G. U. Singco, A. M. Young, K. W. Guarini, and M. Ieong, "Three-dimensional integrated circuits," *IBM J. Res. Develop.*, vol. 50, pp. 491–506, Jul. 2006.

[9] Y. Ye, L. Duan, J. Xu, J. Ouyang, M. K. Hung, and Y. Xie, "3-D optical networks-on-chip (NoC) for multiprocessor systems-on-chip (MPSoC)," in *Proc. IEEE Int. Conf. 3-D Syst. Integr.*, Sep. 2009, pp. 1–6.

[10] B. Feero and P. Pande, "Networks-on-chip in a three-dimensional environment: A performance evaluation," *IEEE Trans. Comput.*, vol. 58, no. 1, pp. 32–45, Jan. 2009.

[11] H. Gu and J. Xu, "Design of 3-D optical network on chip," in *Proc. Symp. Photonics Optoelectron.*, Aug. 2009, pp. 1–4.

[12] Y. Ye, J. Xu, X. Wu, W. Zhang, X. Wang, M. Nikdast, Z. Wang, and W. Liu, "System-level modeling and analysis of thermal effects in optical networks-on-chip," *IEEE Trans. Very Large Scale Integr. Syst.*, vol. PP, no. 99, pp. 1–14, Feb. 2012.

[13] M. Briere, B. Girodias, Y. Bouchebaba, G. Nicolescu, F. Mieyeville, F. Gaffiot, and I. O'Connor, "System level assessment of an optical NoC in an MPSoC platform," in *Proc. Design, Autom. Test Eur. Conf. Exhibit.*, 2007, pp. 1–6.

[14] R. Beausoleil, J. Ahn, N. Binkert, A. Davis, D. Fattal, M. Fiorentino, N. Jouppi, M. McLaren, C. Santori, R. Schreiber, S. Spillane, D. Vantrease, and Q. Xu, "A nanophotonic interconnect for high-performance many-core computation," in *Proc. 16th IEEE Symp. High Performance Interconnects*, Aug. 2008, pp. 182–189.

[15] S. Pasricha and N. Dutt, "ORB: An on-chip optical ring bus communication architecture for multi-processor systems-on-chip," in *Proc. Asia South Pacific Conf. Design Autom.*, Mar. 2008, pp. 789–794.

[16] H. Gu, J. Xu, and W. Zhang, "A low-power fat tree-based optical network-on-chip for multiprocessor system-on-chip," in *Proc. Design, Autom. Test Europe Conf. Exhibition*, Apr. 2009, pp. 3–8.

[17] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonesi, "Phastlane: A rapid transit optical routing network," in *Proc. 36th Annu. Int. Symp. Comput. Architecture*, 2009, pp. 441–450.

[18] A. Kodi, R. Morris, A. Louri, and X. Zhang, "On-chip photonic interconnects for scalable multi-core architectures," in *Proc. 3rd ACM/IEEE Int. Symp. Netw.-on-Chip*, May 2009, p. 90.

[19] D. Ding, Y. Zhang, H. Huang, R. Chen, and D. Pan, "O-router: An optical routing framework for low power on-chip silicon nano-photonic integration," in *Proc. 46th ACM/IEEE Des. Autom. Conf.*, Jul. 2009, pp. 264–269.

[20] S. Xiao, M. H. Khan, H. Shen, and M. Qi, "Multiple-channel silicon micro-resonator based filters for WDM applications," *Opt. Express*, vol. 15, no. 12, pp. 7489–7498, 2007.

[21] A. W. Poon, F. Xu, and X. Luo, "Cascaded active silicon microresonator array cross-connect circuits for WDM networks-on-chip," in *Proc. SPIE 6898, Silicon Photonics III*, 2008, p. 689812.

[22] H. Gu, K. H. Mo, J. Xu, and W. Zhang, "A low-power low-cost optical router for optical networks-on-chip in multiprocessor systems-on-chip," in *Proc. IEEE Comput. Soc. Annu. Symp. VLSI*, May 2009, pp. 19–24.

[23] Y. Xie, M. Nikdast, J. Xu, W. Zhang, Q. Li, X. Wu, Y. Ye, X. Wang, and W. Liu, "Crosstalk noise and bit error rate analysis for optical network-on-chip," in *Proc. ACM/IEEE Design Autom. Conf.*, Jun. 2010, pp. 657–660.

[24] A. Poon, X. Luo, F. Xu, and H. Chen, "Cascaded microresonator-based matrix switch for silicon on-chip optical interconnection," *Proc. IEEE*, vol. 97, no. 7, pp. 1216–1238, Jul. 2009.

[25] R. Ji, L. Yang, L. Zhang, Y. Tian, J. Ding, H. Chen, Y. Lu, P. Zhou, and W. Zhu, "Five-port optical router for photonic networks-on-chip," *Opt. Express*, vol. 19, no. 21, pp. 20 258–20 268, Oct. 2011.

[26] Y. Chen, J. Hu, and X. Ling, "De Bruijn graph based 3D network on chip architecture design," in *Proc. Circuits Syst. Int. Conf. Commun.*, Jul. 2009, pp. 986–990.

[27] S. Mogg, N. Chitica, U. Christiansson, R. Schatz, P. Sundgren, C. Asplund, and M. Hammar, "Temperature sensitivity of the threshold current of long-wavelength InGaAs–GaAs VCSELs with large gain-cavity detuning," *IEEE J. Quantum Electron.*, vol. 40, no. 5, pp. 453–462, May 2004.

[28] V. Padgaonkar and A. Arbor, "Thermal effects in silicon based resonant cavity devices," *NNIN REU Research Accomplishments*, Cornell Univ., 2004, pp. 98–99.

[29] A. Syrbu, A. Mereuta, V. Iakovlev, A. Caliman, P. Royo, and E. Kapon, "10 Gbps VCSELs with high single mode output in 1310 nm and 1550 nm wavelength bands," in *Proc. Conf. Opt. Fiber Commun./Nat. Fiber Optic Engineers*, 2008, pp. 1–3.

[30] C. Kromer, G. Sialm, C. Berger, T. Morf, M. Schmatz, F. Ellinger, D. Erni, G.-L. Bona, and H. Jackel, "A 100-mW 4×10 Gb/s transceiver in 80-nm CMOS for high-density optical interconnects," *IEEE J. Solid-State Circuits*, vol. 40, no. 12, pp. 2667–2679, Dec. 2005.

[31] T. Anan, N. Suzuki, K. Yashiki, K. Fukatsu, H. Hatakeyama, T. Akagawa, K. Tokutome, and M. Tsuji, "High-speed 1.1-μm-range InGaAs VCSELs," in *Proc. Optical Fiber Commun./Nat. Fiber Optic Engineers Conf.*, Feb. 2008, pp. 1–3.
[32] J. A. Lott, N. N. Ledentsov, V. A. Shchukin, A. Mutig, S. A. Blokhin, A. M. Nadtochiy, G. Fiol, and D. Bimberg, "850 nm VCSELs for up to 40 Gbit/s short reach data links," in *Proc. CLEO QELS*, May 2010, pp. 1–2.
[33] W. Liu, J. Xu, X. Wu, Y. Ye, X. Wang, W. Zhang, M. Nikdast, and Z. Wang, "A NoC traffic suite based on real applications," in *Proc. IEEE Comput. Soc. Annu. Symp. VLSI*, Jul. 2011, pp. 66–71.
[34] R. Das, S. Eachempati, A. Mishra, V. Narayanan, and C. Das, "Design and evaluation of a hierarchical on-chip interconnect for next-generation cmps," in *Proc. IEEE 15th Int. Symp. High Performance Comput. Architecture.*, Feb. 2009, pp. 175–186.
[35] W. Liu, M. Yuan, X. He, Z. Gu, and X. Liu, "Efficient SAT-Based mapping and scheduling of homogeneous synchronous dataflow graphs for throughput optimization," in *Proc. Real-Time Syst. Symp.*, Nov. 2008, pp. 492–504.
[36] E. Bonetto, L. Chiaraviglio, D. Cuda, G. Gavilanes Castillo, and F. Neri, "Optical technologies can improve the energy efficiency of networks," in *Proc. 35th Eur. Conf. Opt. Commun.*, 2009, pp. 1–4.
[37] I. O'Connor, "Optical solutions for system-level interconnect," in *Proc. Int. Workshop Syst. Level Interconnect Prediction*, 2004, pp. 79–88.
[38] J. Poulton, R. Palmer, A. Fuller, T. Greer, J. Eyles, W. Dally, and M. Horowitz, "A 14-mW 6.25-Gb/s transceiver in 90-nm CMOS," *IEEE J. Solid-State Circuits*, vol. 42, no. 12, pp. 2745–2757, Dec. 2007.
[39] G. Masini, G. Capellini, J. Witzens, and C. Gunn, "A 1550 nm, 10 Gbps monolithic optical receiver in 130 nm CMOS with integrated Ge waveguide photodetector," in *Proc. 4th IEEE Int. Conf. Group Photonics*, Sep. 2007, pp. 1–3.
[40] F. Gan, T. Barwicz, M. Popovic, M. Dahlem, C. Holzwarth, P. Rakich, H. Smith, E. Ippen, and F. Kartner, "Maximizing the thermo-optic tuning range of silicon photonic structures," in *Proc. Photonics Switching*, Aug. 2007, pp. 67–68.
[41] V. Raghunathan, J. Hu, W. N. Ye, J. Michel, and L. C. Kimerling, "Athermal silicon ring resonators," in *Proc. Integr. Photonics Res., Silicon Nanophoton.*, 2010, p. IMC5.

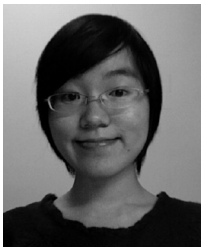

Yaoyao Ye (S'09) received the B.S. degree in electronic engineering from the University of Science and Technology of China, Hefei, China, in 2008. She is currently pursuing the Ph.D. degree in electronic and computer engineering with the Hong Kong University of Science and Technology (HKUST), Kowloon, Hong Kong.

Her current research interests include networks-on-chip, multiprocessor systems-on-a-chip, and embedded systems.

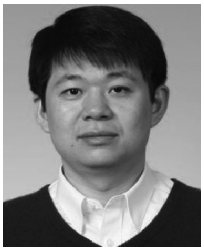

Jiang Xu (S'02–M'07) received the Ph.D. degree from Princeton University, Princeton, NJ, in 2007.

From 2001 to 2002, he was with Bell Laboratories, Murray Hill, NJ as a Research Associate. He was a Research Associate with NEC Laboratories America, Princeton, NJ, from 2003 to 2005. He joined a startup company, Sandbridge Technologies, Tarrytown, NY, from 2005 to 2007 and developed and implemented two generations of NoC-based ultralow power multiprocessor systems-on-chip for mobile platforms. In 2007, he joined the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Clear Water Bay, Kowloon, Hong Kong, as an Assistant Professor, and established the Mobile Computing System Laboratory. He has authored or coauthored more than 60 book chapters and papers in peer-reviewed journals and international conferences. His current research interests include networks-on-chip, multiprocessor systems-on-chip, embedded systems, computer architectures, low-power VLSI design, and HW/SW codesign.

Dr. Xu is currently an Associate Editor of the *ACM Transactions on Embedded Computing Systems* and the IEEE TRANSACTIONS ON VERY LARGE SCALE INTEGRATION SYSTEMS. He is an ACM Distinguished Speaker and a Distinguished Visitor of the IEEE Computer Society. He has served on organizing committees and technical program committees of many international conferences.

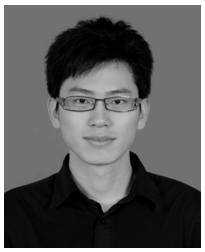

Baihan Huang (S'12) received the B.Eng. degree in electronic and computer engineering from the Hong Kong University of Science and Technology, Kowloon, Hong Kong, in 2011, where he is currently pursuing the M.Phil. degree.

His currrent research interests include brain–machine interface, electroencephalographic signal processing, and pattern recognition.

Xiaowen Wu (S'12) received the B.Sc. degree in computer science from the Harbin Institute of Technology, Harbin, China in 2008. He is currently pursuing the Ph.D. degree in electronic and computer engineering with the Hong Kong University of Science and Technology, Kowloon, Hong Kong.

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

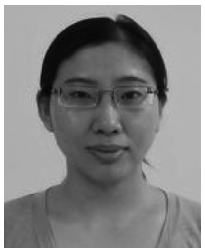

Wei Zhang (M'05) received the B.Eng. and M.Eng. degrees in electrical engineering from the Harbin Institute of Technology, Harbin, China, in 1999 and 2001, respectively and the Ph.D degree from the Department of Electrical Engineering, Princeton University, Princeton, NJ, in 2009.

She joined the School of Computer Engineering, Nanyang Technological University, Singapore, in January 2010. She has published more than 20 papers in international conferences and journals. Her current research interests include embedded systems, systems-on-chip, reconfigurable computing, nanotechnology, and electronic design automation.

Dr. Zhang was a recipient of the Best Paper Award from the IEEE Computer Society Annual Symposium on VLSI in 2009.

Xuan Wang (S'12) received the B.S. degree in electrical engineering from Shanghai Jiaotong University, Shanghai, China in 2009. From 2009, he is currently working toward the Ph.D. degree in the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology (HKUST), Kowloon.

His research interests include embedded system, multiprocessor system, network-on-chip and fault tolerant design, and reliability issues in very deep submicronmeter technologies.

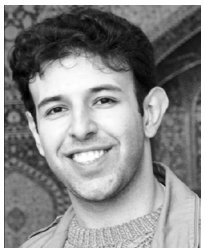

Mahdi Nikdast (S'10) received the B.Sc. degree in computer engineering (Hons.) from Islamic Azad University, Esfahan, Iran, in 2009. Since 2009, he has been pursuing the Ph.D. degree with the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology (HKUST), Kowloon.

He is currently focusing on SNR analyses for ONoCs with the Mobile Computing System Laboratory, HKUST. His current research interests include embedded systems, multiprocessor systems-on-chip, networks-on-chip, and computer architectures.

Mr. Nikdast was a recipient of the Second Best Project Award from the 6th Annual AMD Technical Forum and Exhibition (AMD-TFE 2010).

Zhehui Wang (S'12) received the B.S. degree in electrical engineering from Fudan University, Shanghai, China, in 2010. He is currently pursuing the Ph.D. degree with the Department of Electronic and Computer Engineering, Hong Kong University of Science and Technology, Kowloon.

His current research interests include embedded systems, multiprocessor systems, networks-on-chip, and floorplan design for networks-on-chip.