

Holistic Comparison of Optical Routers for Chip Multiprocessors

(Invited Paper)

Yaoyao Ye, Xiaowen Wu, Jiang Xu, [†]Wei Zhang, Mahdi Nikdast, Xuan Wang

Department of Electronic and Computer Engineering, The Hong Kong University of Science and Technology

[†]School of Computer Engineering, Nanyang Technological University, Singapore

yeyaoyao@ust.hk, jiang.xu@ust.hk, zhangwei@ntu.edu.sg

Abstract—Network-on-chip (NoC) can improve the performance, power efficiency, and scalability of chip multiprocessors (CMPs). However, traditional NoCs using metallic interconnects consume a significant amount of power to deliver high communication bandwidth required in the near future. Optical NoCs are based on CMOS-compatible optical waveguides and optical routers, and promise significant bandwidth and power advantages. In this work, we review different designs of 5x5 and 4x4 optical routers for mesh or torus-based optical NoCs, and compare them for cost of optical resources and optical power loss. Besides, we use a 8x8 mesh-based optical NoC as a case study and analyze the thermal-induced power overhead while using different optical routers. Results show that the number of switching stages in an optical link directly affects the total optical power loss under thermal variations. By using passive-routing optical routers, the maximum number of switching stages in a XY-routing path is minimized to three, and the thermal-induced power overhead in the optical NoC is less than the matched networks using other routers.

Index Terms—Chip multiprocessor, optical network-on-chip, optical router.

I. INTRODUCTION

With the burgeoning complexity of multiprocessor systems, such as chip multiprocessors (CMPs), tens and even hundreds of processor cores are required to be integrated on a single chip. The performance of chip multiprocessors is determined not only by the performance of individual processors, but also by how efficiently they collaborate with one another. Network-on-chip (NoC) has become an alternative communication architecture for pursuit of high performance and power efficiency [1]–[5]. Supported by advances in fabrication and integration of on-chip CMOS compatible optical components, optical NoCs become an attractive candidate to breakthrough the limitation of metallic interconnects in electronic NoCs. Optical NoCs are based on optical routers, optical waveguides and O/E interfaces. The technological and device aspects of a source-based optical link using heterogeneous integration for on-chip data transport was first presented in [6]. Some currently available technologies such as wavelength-division multiplexing (WDM) [7], [8] can be used in optical NoCs to increase the waveguide bandwidth. Regarding energy efficiency, it is predicted that monolithically integrated optical

components will provide the path to the TB/s I/O data rates with near 1pJ/bit [9].

For on-chip optical routers, microresonators (MRs) are widely used as a wavelength-selective optical switch to perform the switching function. Silicon MRs of small size ($5\mu m$ radius) have been demonstrated. The fabrication is based on silicon waveguide with cross-section $500nm \times 200nm$ and the insertion loss is about $0.5dB$ [7]. It was indicated that the DC power consumption of a $12\mu m$ -diameter MR is on the order of $20\mu W$ [8]. Based on the switching function of MRs, several 5x5 optical routers have been proposed for routing in mesh or torus-based optical NoCs [10]–[16]. Cygnus is a low-power non-blocking 5x5 optical router [10]. It uses 16 MRs, six optical waveguides and two optical terminators to implement a 5x5 switching function. Since the waveguides intersect with each other, there are 13 waveguide crossings inside the Cygnus switching fabric. Waveguide crossings in optical NoCs do not affect the bandwidth, but cause more loss and power consumption during packet transmission. Each waveguide crossing introduces about $0.12dB$ insertion loss to the passing optical signals [8]. Although the loss per crossing is small, a large number of crossings in the optical transmission path may lead to significant power loss. Router architectures should be designed to minimize the number of MRs and waveguide crossings while maintaining necessary switching functions. In this work, we compare different 5x5 optical routers for their cost of optical resources, e.g., the number of MRs, waveguide crossings, and optical terminators. We also study the thermal effect on a 8x8 mesh-based optical NoC while using different optical routers and show the power overhead of the NoC due to thermal variations. Results show that the number of switching stages in an optical link directly affects the total optical power loss under thermal variations. By using passive-routing optical routers, e.g., Cygnus, the maximum number of switching stages in the XY-routing optical mesh is minimized to three.

The rest of the paper is organized as follows. Section II shows the comparison of different 5x5 optical routers for mesh-based optical NoCs. Section III reviews several designs of 4x4 optical routers. In Section IV, we use a 8x8 mesh-based optical NoC as a case study and analyze the thermal-induced power overhead while using different optical routers.

Conclusions are drawn in Section V.

II. 5X5 OPTICAL ROUTERS

Optical routers for optical NoCs are built from two types of basic 1x2 optical switching elements (Figure 1), both of which consist of two optical waveguides and one MR. Light signals can propagate along the waveguide and/or be switched to another direction by the MR. MR has different on-state and off-state resonance wavelengths. The basic elements achieve 1x2 optical switching functions by powering on or off the MR. When powered on, the MR has an on-state resonance wavelength λ_{on} . For an input optical signal with a center wavelength λ_{on} , it would be coupled into the MR and directed to the drop port. If the MR is turned off and has a off-state resonance wavelength λ_{off} , the input optical signal would propagate directly to the through port. Multiple basic switching elements may be combined together to implement predefined switching functions. By turning on/off MRs properly, the injected optical signal can be controlled to propagate from an input port to any output port.

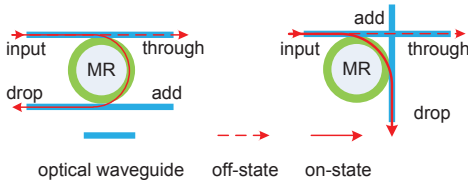


Fig. 1. Basic 1x2 optical switching elements

Figure 2 shows a 5x5 strictly non-blocking optical router Cygnus [10]. The five bidirectional ports include injection/ejection, east, south, west and north ports. They are aligned to their intended directions so no extra crossings will be incurred in the floorplan. Input and output of each port are also properly aligned. The injection/ejection ports are used to connect the processor core through an O/E interface. Cygnus uses 16 MRs to implement a 5x5 switching function. The 5x5 non-blocking crossbar allows five concurrent transactions if there is no contention for the same output port. MRs used in Cygnus are assumed to be resonating at the 1550nm band when they are turned on. The fabrication is based on silicon waveguide with cross-section $500nm \times 200nm$ and the insertion loss is about $0.5dB$ [7]. The MR has a diameter of about $10\mu m$. Optical terminator is an important but expensive device used in the open end of an optical link. Its function is to absorb light and prevent light from returning to the transmission line. In a 5x5 optical crossbar, five horizontal waveguides are crossed with five vertical waveguides. Each waveguide has two ends, one of which is used as input/output and another is open-ended. As a result, ten optical terminators are needed for a 5x5 optical crossbar, with one in each waveguide open end. Cygnus reduces the number of optical terminator to two.

Another important feature of Cygnus is that it can passively route optical signals. By passive routing, we mean that Cygnus

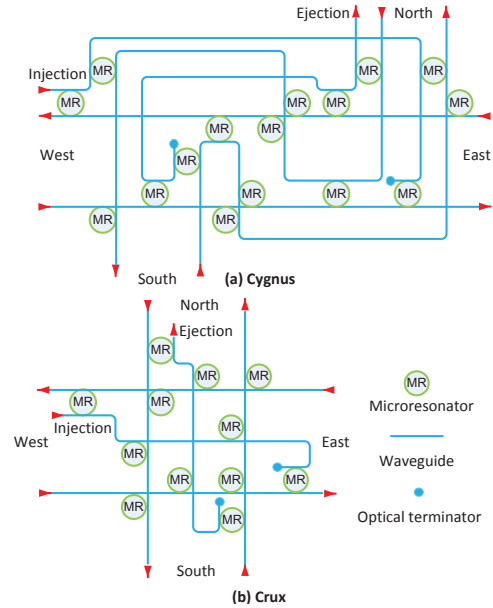


Fig. 2. 5x5 optical router Cygnus [10] and Crux [11]

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 an optical signal turns from one dimension to another dimension 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 optical mesh or torus, including one MR for injection at the source, one MR for a turn during routing, and one MR for ejection at the destination nodes. Based on Cygnus, a more compact 5x5 non-blocking optical router Crux, was proposed [11]. It inherits the passive routing feature of Cygnus. But the switching function of Crux is reduced for XY routing only in mesh or torus-based optical NoCs. XY routing is a low-complexity distributed algorithm without any routing table. Each packet is routed first in X dimension until it reaches the node in the same column with the destination, and then along the perpendicular Y dimension to the destination. They fully utilize the properties of XY routing and reduce the number of MRs in the switching fabric. Compared to Cygnus, Crux uses four less MRs. Besides, Crux takes advantage of the parallel switching element to minimize the waveguide crossing insertion loss. For example, the two waveguides for the injection/ejection port only use the parallel switching elements. Crux reduces the total number of waveguide crossings to nine, and the maximum number of crossings per link from any input port to any output port is five. This feature is beneficial for reducing the total optical power loss, especially for optical NoCs with a large diameter. OXY and ODOR (Figure 3) are two other passive-routing optical routers for XY routing [12], [13].

Besides the optical routers mentioned above, a optimized crossbar-based 5x5 optical router (Figure 4) was demonstrated in [14]. It uses 20 MRs, 10 optical terminators, and 26 wave-

TABLE I
COMPARISON OF DIFFERENT 5X5 OPTICAL ROUTERS

	Cygnus	Crux	OXY	ODOR	[14]	[15]	[16]	[17]
Routing algorithm	Arbitrary	XY routing	XY routing	XY routing	Arbitrary	Arbitrary	Arbitrary	XY routing
Nonblocking	Yes	Yes	Yes	Yes	Yes	Yes	Yes	No
Passive routing	Yes	Yes	Yes	Yes	No	No	No	Yes
N_{MR}	16	12	12	12	20	16	15	12
$N_{Terminator}$	2	2	2	5	10	2	0	2
$N_{Crossing}$	13	9	11	19	26	14	15	12
$L_{average}$	0.78dB	0.64dB	0.73dB	0.87dB	1.15dB	0.87dB	0.77dB	0.96dB

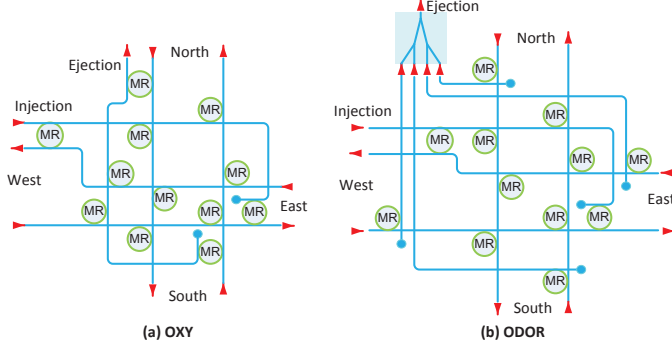


Fig. 3. 5x5 optical router OXY [12] and ODOR [13] for XY routing

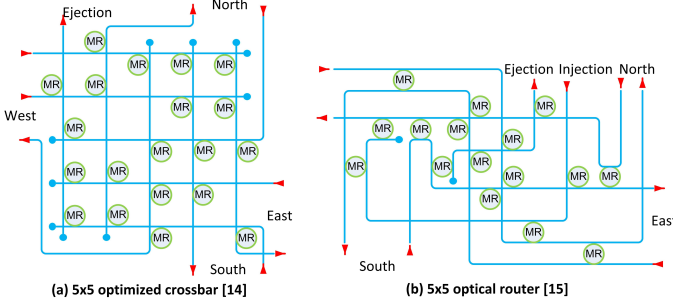


Fig. 4. 5x5 optimized crossbar [14] and router [15]

uide crossings. Different from the passive-routing routers, one MR is required to be turned on in the optimized crossbar for every switching. Ji *et al.* demonstrated another non-blocking 5x5 optical router [15]. The router uses 16 MRs in total, two optical terminators, and 14 waveguide crossings. One MR should be turned on for every switching, except for East-to-South, South-to-East, West-to-North, North-to-West. Min *et al.* proposed a 5x5 optical router which is radially symmetric [16]. It uses 15 MRs and 15 waveguide crossings. Distinguished from other router designs, since it only uses five optical waveguides, no optical terminators are required in this design. A mesh-based optical NoC was proposed by using a 4x4 optical router together with injection gateway and ejection gateway [17]. The 5x5 switching node is blocking though the 4x4 optical router is nonblocking. In this mesh-based optical NoC, at most five MRs will be turned on in any XY-routing optical path, including one MR in injection gateway, one MR

TABLE II
LOSS IN DIFFERENT 5X5 OPTICAL ROUTERS

Loss (dB)	Cygnus	Crux	OXY	ODOR	[14]	[15]	[16]
L_{In_W}	0.50	0.50	0.50	0.98	1.24	1.24	0.50
L_{In_E}	1.00	0.88	1.00	1.12	1.61	0.75	1.00
L_{In_N}	0.75	0.88	0.75	0.99	0.62	0.87	0.75
L_{In_S}	0.75	0.63	1.00	1.12	1.37	1.11	0.75
L_{W_E}	0.51	0.38	0.38	0.50	1.49	1.00	0.75
L_{W_N}	1.00	1.00	1.00	1.24	0.75	0.63	0.50
L_{W_S}	0.50	0.50	0.75	0.87	1.24	0.63	1.00
L_{W_Ej}	1.00	0.88	0.75	0.74	0.62	0.75	0.75
L_{E_W}	0.75	0.38	0.50	0.74	0.87	1.00	1.00
L_{E_N}	0.51	0.50	0.75	0.99	1.24	0.62	0.75
L_{E_S}	1.24	1.00	0.99	1.12	0.74	0.63	0.50
L_{E_Ej}	0.99	0.63	1.00	0.74	1.37	1.00	0.75
L_{N_S}	0.75	0.38	0.50	0.74	0.87	1.00	0.75
L_{N_Ej}	0.74	0.50	0.50	0.50	1.24	0.63	1.00
L_{S_N}	0.51	0.38	0.38	0.62	1.49	1.00	1.00
L_{S_Ej}	1.00	0.88	1.00	0.98	1.61	1.00	0.50

for turning light from south (injection gateway) to X dimension at the source node, one MR for a turn from X dimension to Y dimension in XY routing, one MR for turning light from Y dimension to east (ejection gateway) at the destination node, and one MR in ejection gateway. Table I shows the comparison of their cost of optical resources. By the full 5x5 switching, we mean all possible switchings in a router, except for the U turns. Crux is the most low-cost optical router. It uses 12 MRs and two optical terminators. The total number of waveguide crossing inside Crux is nine, which is less than all other routers listed in the table. Table I shows that Crux has the best loss performance among all the referenced routers. The average loss of passing through Crux in XY routing is about 0.64dB. Table II shows the optical power loss of passing through routers, (e.g. L_{In_W} is the optical power loss from the Injection port to the West output port). We list all possible switchings in XY routing. We assumed the MR on-state insertion loss is 0.5dB, MR off-state passing loss is 0.005dB, the waveguide crossing loss is 0.12dB.

III. 4X4 OPTICAL ROUTERS

Routers at the edge of mesh network do not fully utilize the 5x5 optical switching function. For edge routers, 4x4 optical

TABLE III
COMPARISON OF DIFFERENT 4x4 OPTICAL ROUTERS

	Cygnus-based 4x4	[18]	[19]
Routing algorithm	Arbitrary	Arbitrary	Arbitrary
Nonblocking	Yes	Yes	Yes
Passive routing	Yes	No	No
N_{MR}	8	8	8
$N_{Crossing}$	8	6	10
$L_{average}$	0.66dB	0.58dB	0.74dB

TABLE IV
LOSS IN DIFFERENT 4x4 OPTICAL ROUTERS

Loss (dB)	Cygnus-based 4x4	[18]	[19]
L_{W_E}	0.50	0.75	0.99
L_{W_N}	0.74	0.38	0.62
L_{W_S}	0.75	0.62	0.62
L_{E_W}	0.50	0.75	0.99
L_{E_N}	0.75	0.62	0.62
L_{E_S}	0.74	0.38	0.62
L_{N_W}	0.50	0.38	0.62
L_{N_E}	0.99	0.62	0.62
L_{N_S}	0.50	0.75	0.99
L_{S_W}	0.99	0.62	0.62
L_{S_E}	0.50	0.38	0.62
L_{S_N}	0.50	0.75	0.99

switching functions are enough to satisfy the requirement to connect three neighboring routers and the local O/E interface. In order to further reduce network resource, we can use 4x4 optical switching fabric specifically for the edge routers. Figure 5 shows a Cygnus-based 4x4 optical router, and two other designs from [18] and [19]. The Cygnus-based 4x4 optical router is reduced from the 5x5 Cygnus and inherits the passive routing feature. It does not need to power on any MR for transmissions in the same dimension, i.e. from south to north, east to west and vice-versa. The 4x4 optical router [18] also uses eight MRs as the Cygnus-based 4x4 optical router, and uses two less waveguide crossings. Similar to the 5x5 router in [15], one MR is required to be turned on for every switching, except for East-to-South, South-to-East, West-to-North, North-to-West. The 4x4 optical router [19] uses eight MRs and 10 waveguide crossings. One MR is turned on for every switching, except for East-to-North, North-to-East, West-to-South, South-to-West. Table III shows the comparison of their cost of optical resources. The 4x4 router [18] has a smaller average power loss than other two routers.

IV. THERMAL EFFECT ON A 8x8 MESH-BASED OPTICAL NOC

Thermal sensitivity is a potential issue in optical NoC designs. As a result of thermo-optic effect, the temperature-dependent wavelength shifts in VCSEL (vertical cavity surface emitting laser) and silicon-based MR are found to be about $50\text{-}100\text{pm}/^\circ\text{C}$ [20]. The thermal related wavelength variations will result in additional optical power loss. To ensure that

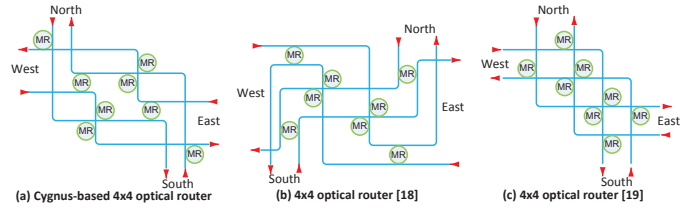


Fig. 5. Cygnus-based 4x4 optical router, router [18], router [19]

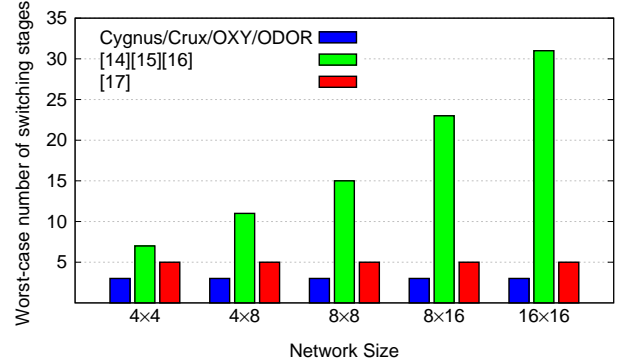


Fig. 6. Worst-case number of switching stages in a mesh-based optical NoC

optical NoCs function properly, a necessary condition is that the optical signal power received by a receiver should be no lower than the receiver sensitivity. As a result, there would be power overhead under thermal variations. Besides, VCSEL power efficiency degrades at high temperatures. Traditional techniques have been proposed to compensate the temperature-dependent wavelength shift for MRs, including active thermal tuning with local microheaters [21]. Based on the system-level analytical optical NoC thermal model and a new passive temperature compensation technique using the optimal device setting proposed in [22], we quantitatively study the thermal effect on an 8x8 mesh-based optical NoC and show the thermal-induced power overhead while using different optical routers.

The number of switching stages in an optical link directly affects the total optical power loss under thermal variations [22]. By using passive-routing optical routers such as Cygnus, Crux, OXY and ODOR, the maximum number of switching stages in a XY-routing path is three regardless of network size. On the other hand, if using crossbar-based optical router [14] or [15], the optical mesh would have a much larger number of switching stages that is proportional to the network size. The worst-case number of switching stages in the 8x8 optical mesh NoC is as large as 15, and the average number of switching stages in a path under uniform traffic is about five. Figure 6 shows the worst-case number of switching stages in a mesh-based optical NoC using different optical routers. If using optical router [14] or [15] or [16], the worst-case number of switching stages in a MxN mesh-based optical NoC would be $M+N-1$. The worst-case number of switching stages in the mesh-based optical NoC [17] is five at most regardless of network size.

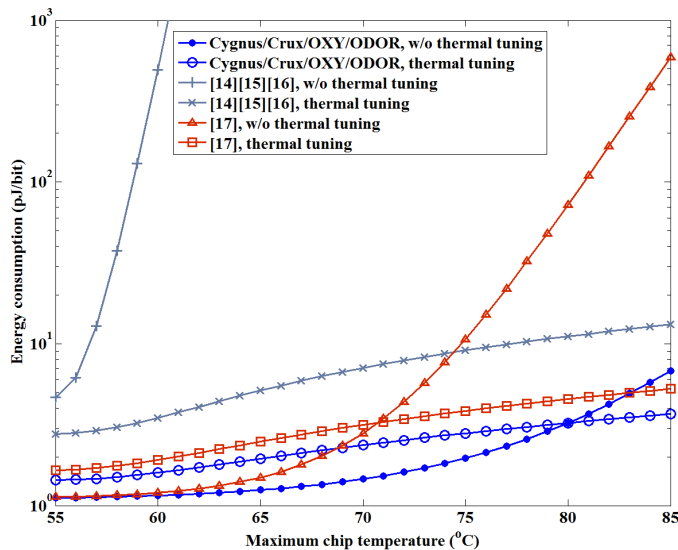


Fig. 7. Worst-case power overhead due to thermal variations

Figure 7 shows the worst-case thermal-induced power overhead in the 8x8 mesh under varying maximum chip temperatures. We assume that the minimum chip temperature is 55°C , and the 3-dB bandwidth of MRs is 3.1nm . Figure 7 shows that by using different optical routers, the worst-case power overhead under thermal variations are different. Optical NoCs that using passive-routing routers have better thermal-induced power efficiency, since the number of switching stages in these NoCs are smaller. When the maximum chip temperature reaches 85°C , with local thermal tuning, the worst-case thermal-induced power overhead is about 3.7pJ/bit if using passive-routing optical router Cygnus (or Crux, OXY, ODOR). If using optical routers [14]–[16], the worst-case thermal-induced power overhead is about 13pJ/bit . If using optical router [17], the worst-case thermal-induced power overhead is about 5.2pJ/bit .

V. CONCLUSIONS

We reviewed different designs of 5x5 and 4x4 optical routers for mesh or torus-based optical NoCs, and compared them for cost of optical resources and optical power loss. Besides, we use a 8x8 mesh-based optical NoC as a case study and analyzed the thermal-induced power overhead while using different optical routers. Results show that the number of switching stages in an optical link directly affects the total optical power loss under temperature variations. By using passive-routing optical routers, the thermal-induced power overhead in the optical mesh NoC is less than by using other routers.

REFERENCES

[1] S. Kumar, A. Jantsch, J.-P. Soininen, M. Forsell, M. Millberg, J. Oberg, K. Tiensyrja, and A. Hemani, "A network on chip architecture and design methodology," in *VLSI Proceedings. IEEE Computer Society Annual Symposium on*, 2002, pp. 105–112.

[2] J. Xu and W. Wolf, "Platform-based design and the first generation dilemma," in *Proceedings of Electronic Design Processes Workshop, Monterey, CA*, April 2002.

[3] W. Dally and B. Towles, "Route packets, not wires: on-chip interconnection networks," in *Design Automation Conference. Proceedings*, 2001, pp. 684–689.

[4] H. G. Lee, N. Chang, U. Y. Ogras, and R. Marculescu, "On-chip communication architecture exploration: A quantitative evaluation of point-to-point, bus, and network-on-chip approaches," *ACM Trans. Des. Autom. Electron. Syst.*, vol. 12, no. 3, pp. 1–20, 2007.

[5] L. Benini and G. De Micheli, "Networks on chips: a new soc paradigm," *IEEE Computer*, vol. 35, no. 1, pp. 70–78, Jan. 2002.

[6] I. O'Connor, F. Tissafi-Drissi, D. Navarro, F. Mieleve, F. Gaffiot, J. Dambre, M. de Wilde, D. Stroobandt, and M. Briere, "Integrated optical interconnect for on-chip data transport," Jun. 2006, pp. 209–209.

[7] 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.

[8] A. W. Poon, F. Xu, and X. Luo, "Cascaded active silicon microresonator array cross-connect circuits for wdm networks-on-chip," J. A. Kubby and G. T. Reed, Eds., vol. 6898, no. 1. SPIE, 2008, p. 689812.

[9] I. Young, E. Mohammed, J. Liao, A. Kern, S. Palermo, B. Block, M. Reshotko, and P. Chang, "Optical i/o technology for tera-scale computing," in *Solid-State Circuits Conference. IEEE International*, 2009, pp. 468–469.

[10] 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 *VLSI. IEEE Computer Society Annual Symposium on*, 13-15 2009, pp. 19–24.

[11] 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 *ACM/IEEE Design Automation Conf.*, June 2010, pp. 657–660.

[12] H. Gu, J. Xu, and Z. Wang, "A novel optical mesh network-on-chip for gigascale systems-on-chip," in *IEEE Asia Pacific Conf. Circuits and Systems*, Nov. 2008, pp. 1728–1731.

[13] H. Gu, J. Xu, and Z. Wang, "Odor: a microresonator-based high-performance low-cost router for optical networks-on-chip," in *Proceedings of the 6th IEEE/ACM/IFIP international conference on Hardware/Software codesign and system synthesis*, 2008, pp. 203–208.

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

[15] 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. 20258–20268, Oct 2011.

[16] R. Min, R. Ji, L. Yang, L. Zhang, Y. Tian, J. Ding, H. Chen, Y. Lu, P. Zhou, and W. Zhu, "Scalable non-blocking optical routers for photonic networks-on-chip," in *IEEE Optical Interconnects Conference*, 2012.

[17] J. Chan, G. Hendry, K. Bergman, and L. Carloni, "Physical-layer modeling and system-level design of chip-scale photonic interconnection networks," *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on*, vol. 30, no. 10, pp. 1507–1520, 2011.

[18] R. Ji, L. Yang, L. Zhang, Y. Tian, J. Ding, H. Chen, Y. Lu, P. Zhou, and W. Zhu, "Microring-resonator-based four-port optical router for photonic networks-on-chip," *Opt. Express*, vol. 19, no. 20, pp. 18945–18955, Sep 2011.

[19] N. Sherwood-Droz, H. Wang, L. Chen, B. G. Lee, A. Biberman, K. Bergman, and M. Lipson, "Optical 4x4 hitless silicon router for optical networks-on-chip (noc): erratum," *Opt. Express*, vol. 16, no. 23, pp. 19395–19395, Nov 2008.

[20] 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 Electronics*, vol. 40, no. 5, pp. 453–462, May. 2004.

[21] 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 *Photonics in Switching*, Aug. 2007, pp. 67–68.

[22] 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," *Very Large Scale Integration (VLSI) Systems, IEEE Transactions on*, vol. PP, no. 99, pp. 1–14, 2012.