

UNION: A Unified Inter/Intra-Chip Optical Network for Chip Multiprocessors

(Invited Paper)

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

[†]Department of Electrical and Computer Engineering, Hong Kong University of Science and Technology, Hong Kong, PRC

[‡]School of Computer Engineering, Nanyang Technological University, 50 Nanyang Avenue, Singapore

[†]Email: {wxxaf, yeyaoyao, weichen, mnkdast, eexwang, eexu}@ust.hk, [‡]zhangwei@ntu.edu.sg

Abstract—As modern computing systems become increasingly complex, communication efficiency among and inside chips has become as important as the computation speeds of individual processor cores. Traditionally, inter-chip and intra-chip communication architectures are separately designed to maximize design flexibility under different constraints. However, jointly designing communication architectures for both inter-chip and intra-chip communication could potentially yield better solutions. In this paper, we present a unified inter/intra-chip optical network, called UNION, for chip multiprocessors (CMP). UNION is based on recent progress in nano-photonic technologies. It connects not only processors on a single CMP but also multiple CMPs in a system. UNION employs a hierarchical optical network to separate inter-chip communication traffic from intra-chip communication traffic. It fully utilizes a single optical network to transmit both payload packets and control packets. The network controller on each CMP not only manages intra-chip communications but also collaborate with each other to facilitate inter-chip communications. We compared CMPs using UNION with those using a matched electronic counterpart in 45 nm process. Based on eight applications, simulation results show that on average UNION improves CMP performance by 3.1X while reducing 92% of network energy consumption and 52% of communication delay.

I. INTRODUCTION

Modern computing systems become increasingly complex to satisfy the growing performance demanded by applications. As the number of transistors available on a single chip increases to billions or even larger numbers, chip multiprocessor (CMP) is becoming an attractive platform for high-performance and low-power applications. In a complex CMP system, the communication efficiency among and inside chips is as important as the computation efficiency of individual processors in the system.

Traditionally, inter-chip and intra-chip communication architectures are separately designed. Intra-chip communication architectures have gradually moved from ad-hoc and bus-based architectures to network-on-chip (NoC) to alleviate the poor scalability, limited bandwidth, and high power consumption of the traditional interconnection networks [1], [2]. As semiconductor technologies continually scale feature sizes down and new applications require even more on-chip communications, conventional metallic interconnects are becoming the bottleneck of NoC. Optical interconnects have been proposed to replace long electrical interconnects in NoC. [3] proposed to

use an optical bus to replace electrical interconnects. [4] proposed Corona to provide high throughput using wavelength-division multiplexing (WDM). [5] proposed an optical NoC, called λ -router, and used WDM technology. [6] proposed a photonic NoC with the topology and routing algorithm. [7] proposed a hybrid optical NoC. [8] proposed a fattree-based optical NoC and integrated the control and data networks. [9] proposed a hybrid mesh-based optical NoC. With steady increasing of individual chip performance, the communications among chips are also blooming. Inter-chip communications still use bus-based and ad-hoc architectures, and signals are transmitted by electrical interconnects on most printed circuit boards (PCB). The limitations of electrical interconnect are already shown in high-performance systems, and optical interconnects are proposed as an alternative to electrical interconnects on PCB [10]. Board-level optical interconnects can use on-board polymer waveguides [11], optical fibers [12] and free space [13] as medium. [14] demonstrated a 160 Gbps chip-to-chip optical data bus using on-board waveguides. [15] proposed an optical processor-to-DRAM network.

Separately designing inter-chip and intra-chip communication architectures can maximize design flexibility under different on-chip and on-board constraints. However, jointly designing communication architectures for both inter-chip and intra-chip communication could potentially yield better solutions. In this paper, we propose a *unified* inter/intra-chip optical network, called UNION. UNION uses nanophotonic technologies to support CMPs. In UNION, data can not only be transmitted optically among processor cores on the same chip, but also be seamlessly transmitted among cores on different chips in optical domain. A collaborative control mechanism is implemented in UNION to facilitate the communications both inside and among chips to improve system performance, delay, and power efficiency.

The rest of paper is organized as follows. Section II details the design of UNION, including its architecture and protocols. Section III compares UNION with a matched electronic network in terms of the performance and energy consumption based on a set of applications. Section IV concludes this paper.

II. UNION ARCHITECTURE

Figure 1 shows an overview of the UNION architecture. UNION includes an inter-chip optical network and intra-

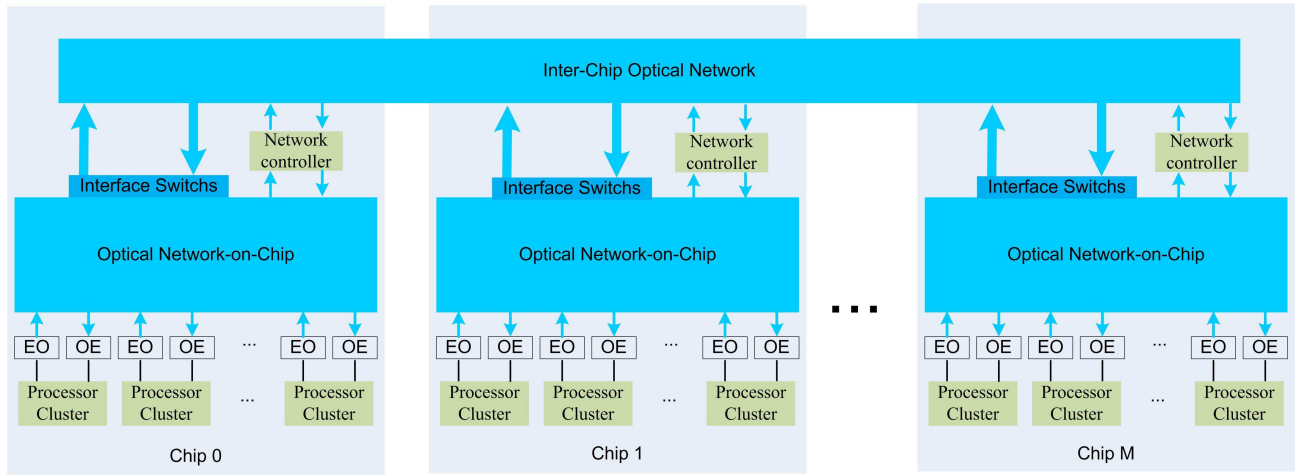


Fig. 1. UNION architecture overview

chip optical networks based on optical NoCs. While intra-chip communications are handled by optical NoCs, inter-chip communications require the collaboration of multiple optical NoCs on different chips through the inter-chip network. Optical NoCs are optically connected to the inter-chip network through interface switches. Each chip has a network controller. The network controllers not only manage the intra-chip networks but also collaborate with each other to facilitate inter-chip communications, which requires both inter-chip and intra-chip networks. In UNION, long electrical interconnects are completely avoided, and there is no optical-to-electrical (OE) or electrical-to-optical (EO) conversions in the middle of paths. In the following, we will detail the intra-chip optical network and inter-chip optical network along with the network protocols.

A. Intra-Chip Network

UNION uses a hierarchical optical NoC for the intra-chip network (Figure 2). The on-chip optical routers in the hierarchical optical NoC are connected in fattree topology. In the fattree topology, each router connects two parent routers via upward links and two children routers through downward links. The top level routers are connected to the inter-chip optical network by interface switches, and the leaf routers are connected to processor clusters by the OE and EO interfaces in concentrators. A processor cluster includes four processor cores and uses an electrical crossbar in the concentrators to communicate. This hybrid approach takes power and performance advantage of short-range electronic network and long-range optical network. All the optical routers are grouped into router clusters and configured by a *network controller* which resides at the top level of fattree. Since the optical loss of each path is different, UNION adjusts the output laser powers in OE interfaces for different optical paths.

1) *Routing Protocol*: In UNION, if both sides of a transaction are within the same concentrator, packets are transmitted totally in electronic domain through a crossbar. On the other hand, if a packet needs to be transmitted out of the

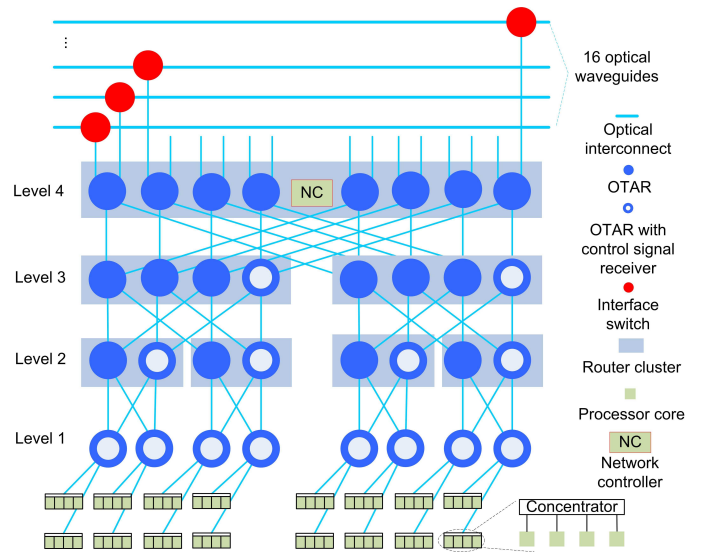


Fig. 2. Intra-chip optical network

concentrator, it first tries to reserve an optical path to the destination concentrator. If the path is reserved successfully, the packet would be transmitted optically to the destination and the destination concentrator would finally switch it to the right core through the local crossbar.

In traditional optical circuit switching, a separate electronic network is needed for path maintenance [6], or the *control* packets can be sent in optical domain but with extra EO/OE conversions at each router along the path [9]. Our design is different from above methods. We implement a special central control unit called *network controller* to configure *all* routers. Especially, all concentrators and routers are *optically* interconnected, and those optical links are neatly combined into a single network. Besides control signal transceivers at both ends of link, no extra components are required.

Network controller contains a buffer storing the states of routers and links. It is responsible for requests arbitration and

path configuration. If a concentrator has data to send, it would send a request with destination information to network controller. After receiving the request, network controller would first find a path based on the routing algorithm detailed in next subsection, and then check all the states of routers and links on the path. If the path is available, network controller would reserve the path and also send a grant signal back to the source core. Failed requests will stay in the network controller until the path is available. Once the source concentrator receives a grant signal, it can send out data propagating along the reserved path. After transmission is finished, a tear down signal will be sent from the core to the network controller to ask for path release. As we can see, only a limited number of control signals need to be transmitted. And compared with distributed path setup mechanisms, UNION can significantly reduce the collisions. These would help to improve the network performance and power efficiency. In following subsections, the routing algorithm and design of our network will be detailed to support this protocol.

2) *Routing Algorithm: Turnaround* routing algorithm is adopted in our fattree network. Specifically, a packet is routed upwards from the source core until it reaches a router which is also the ancestor of the destination core. It is then routed down to the destination. In our implementation, the path is only determined by source/destination information to further ease the network controller. In order to balance the network link utilization, we use shuffling technology to find the path like in [16]. Formally, each router in the upward path in level i checks the packet destination. If the $(i - 1)_{th}$ bit of the destination is 0, we select the left path, otherwise the right path. The downward path is then fixed automatically because of the property of fattree. Network controller chooses the path based on this routing algorithm, and configure the routers for data transmission.

3) *Optical Router:* Optical routers are based on two basic 1×2 switching elements, including the parallel and crossing types. As shown in Figure 3, both of the two switching elements consist of two waveguides and one microresonator (MR). The resonance wavelength of an MR can be controlled by electrical voltage. When the wavelength of input light is the same as the resonance wavelength of MR, the light would be diverted to another waveguide and propagate to the drop port. Otherwise, it would propagate directly to the through port. There may be different MRs with different resonance wavelengths, and each kind of MR can control corresponding light signals while not affecting light in other wavelengths. UNION transmits payload data signals and control signals in wavelengths λ_0 and λ_1 separately.

Based on the two basic switching elements, we can build an optical router, called optical turnaround router (OTAR), for the fattree-based intra-chip optical network. Routers are grouped into *router clusters*, and each cluster as a whole is controlled by an electronic control unit. All clusters are shown in Figure 2, and a Level-2 cluster consisted of two routers is shown in Figure 4.

In Figure 4, the switching fabric of each OTAR router in




Fig. 3. Two basic switching elements

a cluster implements a 4×4 switching function for optical data signals in wavelength λ_0 . The routing functions can be achieved by turning on/off corresponding MRs. It is designed to minimize the number of waveguide crossings. Based on the routing algorithm, some turns in the router can be eliminated. Specifically, there is no U-turn and turns between up-left and up-right ports. One of the routers of a cluster is different from the others. The right router is attached with a control signal receiver. The MR with resonance wavelength λ_1 will direct the control signals from network controller to a router control unit. The received control information would be interpreted to configure *all* MRs in this cluster with resonance wavelength λ_0 . After MRs configuration, path is setup for payload data signals in wavelength λ_0 . Top level routers are also attached with MRs in resonance wavelength λ_1 , responsible for receiving control data from source concentrators and sending out control data to destination concentrators and clusters.




Fig. 4. A Level-2 router cluster including two OTARs

With the above designs of router and clusters, all upward paths from cores to the network controller and downward paths from the network controller to clusters, are distinct without any overlap. As a result, network controller can connect all clusters and concentrators in a point-to-point fashion. And a single optical network is used for both data and control information. In following section, we would show how the inter-chip network is designed and how it is connected to the intra-chip network.

B. Inter-Chip Network

The inter-chip network connects all the intra-chip networks. In UNION, we designed an optical bus with distributed control for inter-chip network (Figure 5). Network controllers

collaboratively arbitrate the optical bus and manage their own intra-chip network resource for inter-chip communications. Although bus-based communication architectures have limited scalability, they are still a viable low-cost choice for systems with a moderate number of chips. UNION’s inter-chip network consists of an optical data bus (at top of Figure 5) and an optical control bus (at bottom of Figure 5). The data bus is responsible for data communications between chips, and the control bus helps network controllers to cooperate with each other during bus arbitration.

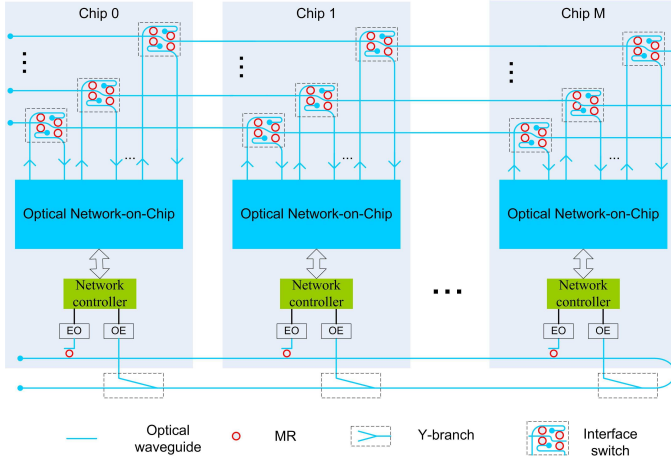


Fig. 5. Inter-chip optical network

1) *Optical Data Bus*: In UNION’s inter-chip network, the number of data bus channels is proportional to the number of top level routers in the intra-chip network. Each data bus channel is composed of a on-chip silicon waveguide, a polymer waveguide embedded on PCB board, and optical connectors which connect on-chip waveguides with on-board waveguides. Each channel is bidirectional and half-duplex. For 64-core CMPs, only 16 data bus channels are required. We designed *interface switches* to connect top level routers in the intra-chip network to optical data bus channels, as shown at top of Figure 5. The interface switch is composed of four MRs and two waveguides. Data signals can be sent to the bus in either direction depending on which MR is powered on. Insertion loss caused by interface switch is minimized in the design. If no MR is powered on, data signals will pass current chip with little optical power loss.

A useful feature of our optical data bus design is that a single data channel can be used by multiple chips simultaneously. Interface switches can divide a single data channel into multiple sections using the unidirectional property of optical signals, and each section can operate independently. The distributed arbitration can utilize this feature to reduce data collisions and improve performance.

2) *Optical Control Bus*: Since multiple chips can send data out simultaneously, arbitration is required to avoid collisions. The bus arbitration is made *collaboratively* by the network controllers. A control bus is implemented to help them cooperate with each other, shown at the bottom of Figure 5. The

control bus is primarily a waveguide which connects all the network controllers. It allows a network controller to *broadcast* control signals. As shown in the figure, an MR is used to inject control signals into the control bus, and a Y-branch is used to eject control signals. Y-branches are designed with different split ratio. The $(N - i)$ th Y-branch from left to right has a split ratio of $i : 1$, and this allows the network controllers to receive the same amount of power.

3) *Network Protocols*: Inter-chip communications require both the intra-chip and inter-chip networks, and are managed collaboratively by the network controllers. When a processor core wants to start a communication with another core on a different chip, it first sends a request to the network controller through a concentrator, which is the same as an *intra-chip* communication. After receiving the request, the network controller will broadcast it to the network controller on the destination chip. The source and destination network controllers will simultaneously start to reserve an on-chip up-link path and down-link path respectively. They will use the same deterministic routing algorithm as for intra-chip communications. Network controllers will broadcast successful path reservations on the control bus. When both the on-chip up-link and down-link paths are reserved, network controllers will reserve a data bus channel and sends a grand signal to the source processor. After receiving the grand signal, source processor will send immediately. Upon finishing the data transmission, a tear down signal is sent from the source core to the source network controller, which in turn broadcast it to the destination controller. All network controllers will update their status buffers based on received information.

III. EVALUATION AND RESULT

We compared UNION with a matched electronic network composed of a fattree-based electronic NoC and inter-chip bus in terms of performance, energy consumption and delay. Eight applications are used for the comparison, including H263 encoder, H263 decoder, satellite receiver, sample rate converter, modem, and H264 decoder with different rates. For each application, an offline optimization approach is applied for mapping and scheduling tasks onto CMPs with the objective of maximizing system performance. We developed SystemC-based cycle-accurate simulators for UNION and its counterpart. We simulated both the networks for eight chips and each chip is a 64-core CMP.

A. Performance Comparison

Performance is measured in terms of the average number of iterations that an application can finish in a given time. In the electronic fattree NoC, the same turnaround routing algorithm is employed for packet switching. Wormhole routing is adopted to avoid head-of-line (HOL) problem and improve performance, and back pressure is used for flow control. The electronic routers are pipelined, and virtual channels are implemented. We assumed the routers are running at 1.25 GHz, and each port is 32-bit wide and bidirectional. Each 32-bit flit can be transmitted in one clock cycle, and the link

bandwidth is 40 Gbps. For the electrical inter-chip bus, we assumed that each link works at 10 Gbps [17]. There are 64 bidirectional links which connects 32 top level routers, and thus the bisectional bandwidth of the bus is 640 Gbps.

In UNION, we assumed that electronic components are also running at 1.25 GHz. For comparison, we also assumed the link bandwidth is 40 Gbps. Every four cores are connected to an electronic concentrator. The 16 concentrators are connected with the intra-chip optical network. Therefore, the bisectional bandwidth of the UNION intra-chip network is only a quarter of the electronic NoC. There are 16 bidirectional data bus channels given the same bisectional bandwidth as the electrical bus. We implemented the network controller in VHDL and synthesized it with a 45nm library. The network controller can simultaneously handle 16 requests in 20 clock cycles based on the synthesis result.




Fig. 6. Normalized performance of UNION compared to the electronic counterpart for different applications

Figure 6 shows the normalized performance of each application on CMPs using UNION compared to the electronic counterpart. For most of applications, CMPs using UNION achieve more than 3X improvement compared with the CMPs using the electronic counterpart. Satellite receiver application only shows 1.1X improvement because the application’s data flow is mostly confined in individual CMPs. On average, UNION help to improve the CMP performance by 3.1X. Figure 7 shows the normalized average end-to-end (ETE) delays of the applications in UNION compared with the match electronic network. On average, the ETE communication delay of UNION is only 48% of its electronic counterpart. The satellite receiver application also shows less improvement. While considering UNION’s intra-chip network has only 25% bisectional bandwidth of the matched electronic network, UNION utilizes its network recourses more effectively.

B. Energy Evaluation and Comparison

UNION consumes power in several ways, including payload data power consumption and control power consumption. Payload data power consumption involves the concentrators, MRs in OTAR, and EO and OE interfaces for both intra-chip or inter-chip communications. The EO interfaces include serializer/deserializer [18], VCSEL [19] and driver [20].




Fig. 7. Normalized ETE delay in UNION compared to the electronic counterpart for different applications

The OE interfaces include photodetector [21] and the TIA-LA circuits [20]. Optical power loss dominates the power consumption of the system. It can be estimated based on the loss of each optical component. The MR insertion loss is 0.5 dB. The silicon waveguide crossing insertion loss, MR passing loss, waveguide bending loss and waveguide propagation loss are 0.12 dB, 0.005 dB, 0.005 dB/90° and 0.17 dB/mm respectively [22] [23] [24]. The coupling loss between on-chip and on-board waveguides is 0.45 dB [25]. The propagation loss on the polymer waveguide on PCB is 0.035 dB/cm [26]. As for the electronic network, the electronic router and metal wires were simulated in Cadence Spectre, and power characteristics were derived based on the simulations. For the interconnect power consumption in the electronic bus, we used the latest result from [17].




Fig. 8. Normalized energy consumption of UNION compared to the electronic counterpart for different applications

Figure 8 shows the normalized energy consumption of UNION compared to the electronic counterpart for different applications. On average, UNION consumes 92% less energy compared with the matched electronic network. Satellite receiver application has the lowest improvement of 80%. Further analysis shows that the high energy efficiency is in both intra-chip and inter-chip communications. In the electronic NoC, long metallic interconnects and buffers consume large amount

of power to delivery required bandwidth. Optical interconnects in UNION significantly lower the energy consumption, and optical signals are transmitted from source to destination without buffering. As for inter-chip communications, the optical bus also consumes significantly lower energy than the electronic bus. The adaptive power control mechanism further improves UNION's energy efficiency.

IV. CONCLUSION

A unified inter/intra-chip optical interconnection network, called UNION, for CMPs is proposed in this paper. We jointly designed the inter-chip and intra-chip networks in UNION. UNION employs a hierarchical optical network to separate inter-chip communication traffic from intra-chip communication traffic. It fully utilizes a single optical network to transmit both payload packets and control packets. The network controller on each CMP not only manages intra-chip communications but also collaborate with each other to facilitate inter-chip communications. We compared CMPs using UNION with those using a matched electronic counterpart in 45 nm process. Based on eight applications, simulation results show that on average UNION improves CMP performance by 3.1X while reducing 92% of network energy consumption and 52% of communication delay.

ACKNOWLEDGMENT

This work is partially supported by RGC of the Hong Kong Special Administrative Region, China.

REFERENCES

- [1] W. Dally and B. Towles, "Route packets, not wires: on-chip interconnection networks," in *Design Automation Conference, 2001. Proceedings*, 2001, pp. 684–689.
- [2] J. Xu, W. Wolf, J. Henkel, and S. Chakradhar, "A design methodology for application-specific networks-on-chip," *ACM Trans. Embed. Comput. Syst.*, vol. 5, no. 2, pp. 263–280, 2006.
- [3] N. Kirman, M. Kirman, R. K. Dokania, J. F. Martinez, A. B. Apsel, M. A. Watkins, and D. H. Albonesi, "Leveraging optical technology in future bus-based chip multiprocessors," in *MICRO 39: Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*. Washington, DC, USA: IEEE Computer Society, 2006, pp. 492–503.
- [4] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn, "Corona: System implications of emerging nanophotonic technology," in *ISCA '08: Proceedings of the 35th International Symposium on Computer Architecture*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 153–164.
- [5] I. O'Connor, F. Mieyeville, F. Gaffiot, A. Scandurra, and G. Nicolescu, "Reduction methods for adapting optical network on chip topologies to specific routing applications," in *Proc. Design of Circuits and Integrated Systems (DCIS)*, November 12–14 2008.
- [6] A. Shacham, K. Bergman, and L. P. Carloni, "Photonic networks-on-chip for future generations of chip multiprocessors," *IEEE Trans. Comput.*, vol. 57, no. 9, pp. 1246–1260, 2008.
- [7] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonesi, "Phastlane: a rapid transit optical routing network," in *ISCA*, 2009, pp. 441–450.
- [8] K. H. Mo, Y. Ye, X. Wu, W. Zhang, W. Liu, and J. Xu, "A hierarchical hybrid optical-electronic network-on-chip," in *IEEE Computer Society Annual Symposium on VLSI*, 2010.
- [9] H. Gu, J. Xu, and W. Zhang, "A low-power fat tree-based optical network-on-chip for multiprocessor system-on-chip," in *DATE*, 2009, pp. 3–8.
- [10] D. A. B. Miller, "Physical reasons for optical interconnection," *Special Issue on Smart Pixels, International Journal of Optoelectronics*, vol. 11, no. 3, pp. 155–168, 1997.
- [11] G. Van Steenberge, P. Geerinck, S. Van Put, J. Van Koetsem, H. Otevaere, D. Morlion, H. Thienpont, and P. Van Daele, "MT-compatible laser-ablated interconnections for optical printed circuit boards," *Lightwave Technology, Journal of*, vol. 22, no. 9, pp. 2083–2090, Sept. 2004.
- [12] S. H. Hwang, M. H. Cho, S.-K. Kang, H.-H. Park, H. S. Cho, S.-H. Kim, K.-U. Shin, and S.-W. Ha, "Passively assembled optical interconnection system based on an optical printed-circuit board," *Photonics Technology Letters, IEEE*, vol. 18, no. 5, pp. 652–654, 1, 2006.
- [13] A. Apsel, Z. Fu, and A. Andreou, "A 2.5-mw SOS CMOS optical receiver for chip-to-chip interconnect," *Lightwave Technology, Journal of*, vol. 22, no. 9, pp. 2149–2157, Sept. 2004.
- [14] F. Doany, C. Schow, R. Budd, C. Baks, D. Kuchta, P. Pepeljugin, J. Kash, F. Libsch, R. Dangel, F. Horst, and B. Offrein, "Chip-to-chip board-level optical data buses," in *Optical Fiber Communication/National Fiber Optic Engineers Conference, 2008. OFC/NFOEC 2008. Conference on*, Feb. 2008, pp. 1–3.
- [15] 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 *HOTI '08: Proceedings of the 2008 16th IEEE Symposium on High Performance Interconnects*. Washington, DC, USA: IEEE Computer Society, 2008, pp. 21–30.
- [16] C. G. Requena, F. G. Villamón, M. E. Gómez, P. López, and J. Duato, "Deterministic versus adaptive routing in fat-trees," in *IPDPS*. IEEE, 2007, pp. 1–8.
- [17] G. Balamurugan, J. Kennedy, G. Banerjee, J. Jaussi, M. Mansuri, F. O'Mahony, B. Casper, and R. Mooney, "A scalable 5-15 Gbps, 14-75 mw low-power I/O transceiver in 65 nm CMOS," *Solid-State Circuits, IEEE Journal of*, vol. 43, no. 4, pp. 1010–1019, april 2008.
- [18] 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," *Solid-State Circuits, IEEE Journal of*, vol. 42, no. 12, pp. 2745–2757, dec. 2007.
- [19] A. Syrbu, A. Mereuta, V. Iakovlev, A. Caliman, P. Royo, and E. Kapon, "10 Gbps VCSELs with high single mode output in 1310nm and 1550 nm wavelength bands," in *Optical Fiber Communication/National Fiber Optic Engineers Conference, 2008. OFC/NFOEC 2008. Conference on*, Feb. 2008, pp. 1–3.
- [20] C. Kromer, G. Sialm, C. Berger, T. Morf, M. Schmatz, F. Ellinger, D. Erni, G.-L. Bona, and H. Jackel, "A 100mw 4 × 10Gb/s transceiver in 80nm CMOS for high-density optical interconnects," in *Solid-State Circuits Conference, 2005. Digest of Technical Papers. ISSCC. 2005 IEEE International*, Feb. 2005, pp. 334–602 Vol. 1.
- [21] G. Masini, G. Capellini, J. Witzens, and C. Gunn, "A 1550nm, 10 Gbps monolithic optical receiver in 130nm CMOS with integrated Ge waveguide photodetector," in *Group IV Photonics, 2007 4th IEEE International Conference on*, Sept. 2007, pp. 1–3.
- [22] A. W. Poon, F. Xu, and X. Luo, "Cascaded active silicon microresonator array cross-connect circuits for WDM networks-on-chip," in *Silicon Photonics III*, vol. 6898, no. 1. SPIE, 2008, p. 689812.
- [23] 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.
- [24] F. Xia, L. Sekaric, and Y. Vlasov, "Ultra-compact optical buffers on a silicon chip," *Nature Photonics*, no. 1, pp. 65–71, Jan. 2007.
- [25] J. K. Doyle and A. P. Knights, "Design and simulation of an integrated fiber-to-chip coupler for silicon-on-insulator waveguides," *Selected Topics in Quantum Electronics, IEEE Journal of*, vol. 12, no. 6, pp. 1363–1370, Nov.-Dec. 2006.
- [26] G. L. Bona, B. J. Offrein, U. Bapst, C. Berger, R. Beyeler, R. Budd, R. Dangel, L. Dellmann, and F. Horst, "Characterization of parallel optical-interconnect waveguides integrated on a printed circuit board," in *Micro-Optics, VCSELs, and Photonic Interconnects*, vol. 5453, no. 1. SPIE, 2004, pp. 134–141.