

Photonic NoC for DMA Communications in Chip Multiprocessors

Assaf Shacham, Benjamin G. Lee,
Aleksandr Biberman, Keren Bergman
Columbia University
Dept. of Electrical Engineering
500 W 120th St.,
New York, NY 10027

{assaf,benlee,biberman,bergman}@ee.columbia.edu

Luca P. Carloni
Columbia University
Dept. of Computer Science
1214 Amsterdam Ave.,
New York, NY 10027
luca@cs.columbia.edu

Abstract

*As multicore architectures prevail in modern high-performance processor chip design, the communications bottleneck has begun to penetrate on-chip interconnects. With vastly growing numbers of cores and on-chip computation, a high-bandwidth, low-latency, and, perhaps most importantly, low-power communication infrastructure is critically required for next generation chip multiprocessors. Recent remarkable advances in silicon photonics and the integration of photonic elements with standard CMOS processes suggest the use of **photonic networks-on-chip**. In this paper we review the previously proposed architecture of a hybrid electronic/photonic NoC. We improve the former internally blocking switches by designing a non-blocking photonic switch, and we estimate the optical loss budget and area requirements of a practical NoC implementation based on the new switches. Additionally, we tackle one of the key performance challenges: the latency associated with setting-up photonic paths. Simulations show that the technique suggested can substantially reduce the latency and increase the effective bandwidth. Finally, we consider the DMA communication model in the context of the photonic network and evaluate the optimal DMA block size.*

1. Introduction

The recent emergence of chip multiprocessors (CMPs) for driving performance via increases in the number of parallel computational cores has fundamentally shifted the role of the system interconnects and global communications infrastructure. The future trend is clearly on path toward further multiplication of the on-chip processing cores with perhaps the most pronounced evidence in the recent unveiling of Intel's 80-core multiprocessor that delivers computing

performance exceeding 1 TeraFLOP [18]. As CMP architectures begin to essentially resemble highly parallel computing systems on-chip, their performance is directly tied to how efficiently the parallelism of the system is exploited and its aggregate compute power harnessed. Achieving maximum utilization of compute resources as the number of functional parallel units scales, falls increasingly on the efficiency of the information exchange among those resources. Thus, the global on-chip communications plays a central and dominant role in the ultimate CMP system performance.

The realization of a scalable on-chip communications infrastructure faces critical challenges in meeting the enormous bandwidths, capacities, and stringent latency requirements demanded by CMPs in a power efficient fashion [11, 18]. Recent research on packet-switched networks-on-chip (NoC) [2, 4, 8, 12] has shown that carefully-engineered shared links can provide enough bandwidth to replace many traditional bus-based and/or point-to-point links. However, with a fixed upper limit to the total chip power dissipation, and the communications infrastructure emerging as a major power consumer, performance-per-watt is becoming the most critical design metric for the scaling of NoCs and CMPs. It is not clear how electronic NoCs will continue to satisfy future communication bandwidths and latency requirements within the power dissipation budget.

Over the past several years, integrated photonic technology has witnessed unprecedented advancements in fabrication capabilities of nano-scale devices and precise control over their optical properties. Importantly, these breakthroughs have led to the development of silicon photonic device integration with electronics directly in commercial CMOS manufacturing platforms [6]. Photonic elements are now available as library cells in standard CMOS processes. For the first time, we can consider the practical insertion of photonic interconnection networks with the tremendous bandwidth, latency, and energy advantages that have made

the technology ubiquitous in long-haul transmission systems, as the communications infrastructure for CMPs.

Photonic interconnection networks offer a disruptive technology solution with fundamentally low power dissipation, ultra-high communications bandwidths, and minimal access latencies. For CMPs, photonic NoCs deliver a dramatic reduction in power expended on intrachip global communications. Photonic NoCs essentially change the power scaling rules: as a result of the low loss in optical waveguides, once a photonic path is established, the data are transmitted end-to-end without the need for repeating, regeneration or buffering. In electronic NoCs, on the other hand, a message is buffered, regenerated and then transmitted on the inter-router links multiple times en route to its destination. Furthermore, the switching and regenerating elements in CMOS consume dynamic power that grows with the data rate. The power consumption of optical switching elements, conversely, is independent of the bit rate, so high bandwidth messages do not consume additional dynamic power.

In previous work [16] we provided a detailed power analysis of a photonic NoC, and compared it to an electronic NoC designed to provide the same bandwidth to the same number of cores. The compelling conclusion of the study was that the power expended on intrachip communications can be reduced by **two orders of magnitude** when **high-bandwidth communications** is required among the cores.

Optical technology, however, presents design challenges which are fundamentally different from those faced by electronic NoC designers. While in CMOS technology buffers and processing resources are abundant and are amply used, they are very difficult to implement in optics. Considering the small chip-scale area, no buffers or all-optical processing can be used. To utilize the advantages of photonics in constructing a NoC, we therefore adopt a hybrid approach: a network of silicon optical switches and waveguides is used for bulk message transmission, and an electronic network, with the same topology, is used for distributed control and short message exchange.

In previous work [17] we have presented the network architecture and tackled some design issues such as the construction of the basic switching elements and routers, message sizing, deadlock avoidance, and overcoming the lack of buffering using path overprovisioning. In this paper, we begin with a brief architecture overview, which includes a newly designed non-blocking photonic switch, and an analysis of the optical loss budget and footprint for a specific NoC implementation. Then, we provide a continuation of the previous work's design exploration by re-evaluating the path-setup procedure and message sizing policy. We find that the time spent by the path-setup packets in the internal router buffers is a major contributor to the latency, especially when the network is heavily loaded. By reducing the buffering depth we show that path setup latency is signifi-

cantly reduced, and throughput is improved by as much as 30%.

2 Architecture Overview

Photonic technology offers unique advantages in terms of energy and bandwidth but lacks two necessary functions for packet switching: buffering and processing, which are very difficult to implement. Electronic NoCs, conversely, have many advantages in flexibility, abundant functionality and ample buffering space, but their transmission bandwidth per line is limited. The photonic NoC architecture employs a hybrid design, where an optical interconnection network is used for very-large-bandwidth message transmission, and an electronic network, with the same topology, is used to control the optical network. The topology of both networks is a 2-D torus, a topology that maps well on CMP planar layout. Each core in the CMP is equipped with a network interface, a *gateway*, whose goal is to perform the necessary Electronic/Optical and Optical/Electronic (E/O and O/E) conversions, communicate with the control network and execute several other related tasks (synchronization, etc.)

Every photonic message transmitted is preceded by an electronic control packet (a *path-setup* packet) which is routed on the electronic network, acquiring and setting-up a photonic path for the message. Buffering of messages, which is impossible in the photonic network, only takes place for the electronic packets during the path-setup phase. The photonic messages are transmitted without buffering once the path has been acquired. This approach has many similarities with optical circuit switching, a technique used to establish long lasting connections between nodes in the optical internet core [15].

This section will briefly describe the main issues in the architecture and the design of the hybrid photonic NoC, as discussed in greater detail in [17]. Additionally, a newly designed non-blocking photonic switch, and a feasibility analysis of the optical losses and area consumption resulting from the photonic components in a specific network implementation are included.

Building Blocks. In previous work, the fundamental building block of the photonic network, a broadband *photonic switching element* (PSE) based on microring resonator structures, was introduced. This switch was, in essence, a waveguide intersection, positioned between two ring resonators (Fig. 1). The rings have a certain resonance frequency, which is derived from material and structural properties. Therefore, shifting the resonance frequency provides switching action. The definition of *ON* and *OFF* described below is different than in previous papers. Here, the *ON*




Figure 1. Photonic switching element: (a) ON state: a passive waveguide crossover. (b) OFF state: light is coupled into rings and forced to turn.

and *OFF* states are defined in order to minimize the insertion loss through the rings. In this manner, when the switch is turned *ON* by the injection of electrical current into *p-n* contacts surrounding the rings, the resonances shift, such that the incident light signals are off resonance, and the light passes through the waveguide intersection uninterrupted, as in a passive waveguide crossover (Fig. 1a). In the *OFF* state, the resonance frequency of the rings coincides with the wavelength (or wavelengths) on which the optical data stream is modulated; here, the encoded light is coupled into the rings and onto the perpendicular waveguides, making right-angle turns (Fig. 1b). Thus, switching action is demonstrated by the PSE.

Photonic switching elements and modulators based on microring resonators have been implemented in silicon, and a switching time of 30 ps has been experimentally demonstrated [20]. Their physical dimensions are very small ($5\mu\text{m}$ ring radius), and their power consumption is extremely low: less than 0.5 mW when *ON* and about 1 pJ to switch [20]. When the switches are *OFF*, they act as passive devices and consume nearly no power. They also exhibit good crosstalk properties (> 20 dB), and low insertion loss (approximately 1.5 dB) [19]. These switches are typically narrow-band, but advanced research efforts are now underway which are endeavoring to fabricate wideband structures capable of switching several wavelengths simultaneously, each modulated at tens of Gb/s. It is also reasonable to assume that the loss figures can be improved with advances in fabrication techniques.

In former work [17], an arrangement of the PSEs into a structure creating a 4×4 switch was discussed. In this structure, the PSEs were grouped in fours and interconnected by silicon photonic waveguides. Each quadruplet, controlled by an electronic circuit termed an *electronic router*, formed a 4×4 switch (Fig. 2). The 4×4 switches were interconnected by the inter-PSE waveguides and by metal lines connecting the electronic routers. Control packets (e.g. path-setup) were received in the electronic router, processed, and sent to their next hop, while the PSEs were switched *ON*




Figure 2. 4×4 switch composed of 4 PSEs controlled by an electronic router. Although the layout of the PSEs is later reconfigured for optimal performance, the relationship between the photonic and electronic network elements remains the same.

and *OFF* accordingly. Once a path-setup packet completed its journey through a sequence of electronic routers, a chain of PSEs was ready to route the optical message.

In this work, the same path-setup algorithm is utilized, but the former, simple arrangement of the photonic devices within the routing switches is traded for a design that achieves higher functionality. The previously studied 4×4 switch was internally blocking, requiring specific routing algorithms to improve performance [17]. The new strictly non-blocking switch (Fig. 3a) alleviates this problem by increasing the number of internal paths within the switch. However, the number of microrings remains the same, which is important for power consumption arguments. The redesigned switch guarantees an internal path from any input to any output, as long as no two packets contend for the same output, and as long as packets are not allowed to ingress and egress from the same port (no U-turn).

Owing to the small footprint of the PSEs and the simplicity of the electronic router, which only handles small control packets, the new routing switch can still occupy a very small area. There is a trade-off between the footprint of a microring and the wavelength spacing of the optical signal, which corresponds to the free spectral range of the resonator. The microring diameter is approximately inversely proportional to the wavelength spacing [20]. Based on the size of the microring resonator devices required to implement 1.6-nm wavelength spacing, and the minimal logic required to implement the electronic router, we estimate the total area of the non-blocking routing switch to be $525\mu\text{m} \times 525\mu\text{m}$.

Topology. Two dimensional planar topologies such as meshes and tori are typically chosen as NoC topologies for




Figure 3. Layout of the photonic components for the (a) routing, (b) injection, and (c) ejection switches. The shapes in the upper left corner of each sub-figure signify the role of the switch in the layout of Fig. 4.

CMPs [14]. These topologies are especially suitable considering the small radix of the 4×4 switch. While recent work suggests building fat-tree interconnection networks using high-radix routers [10], these high-radix routers are very difficult to implement with photonics. Tori, which offer a lower network diameter compared to meshes at the expense of having longer links [5], are the obvious choice since the transmission power on photonic links is independent of the length, in contrast to copper lines. We therefore use a folded torus topology for the network. Regular 2-D topologies require 5×5 switches, composed of four I/O ports for network traffic and one for local packet injection/ejection, which are overly complex to implement using photonic technology. The torus is thus augmented with gateway access points (GAPs) connected to the network interfaces in the cores.

The GAPs are designed with two goals in mind: (1) to facilitate injection and ejection without interference with the through traffic on the torus, and (2) to avoid blocking between injected and ejected traffic. These goals are achieved by using 3 types of switches in each GAP: (1) a *gateway switch* is directly connected to the gateway in the processor core; (2) *injection switches* are located on the torus rows; and (3) *ejection switches* are located on the torus columns. Each injected message travels from the gateway switch to an injection switch. It then travels on the network to the ejection switch associated with its destination core from which it is sent to the gateway switch and out of the network. The design of the GAPs and the traffic rules used to avoid message blocking in them are discussed in detail in [17].

A network designer may take advantage of the small footprint to improve the performance by increasing **path multiplicity**. Network contention is a major source of latency in the path setup procedure (Section 3 discusses this in detail). The torus network can be augmented with addi-

tional paths so that the probability of contention is lowered and the path-setup latency is, accordingly, reduced. Owing to the small footprint of the switches, the simplicity of the routers, and the fact that the microrings only consume power when they are activated, the power and area cost of adding parallel paths is not large. Path multiplicity is therefore used as a cost-effective method of improving performance and reducing contention in the absence of traditional means for contention resolution, such as buffers.

In order to demonstrate the feasibility of such a network from the physical layer, we estimate the optical losses and area resulting from a 36-node (6×6) implementation of a 2-D folded torus with $2 \times$ path over-provisioning. A 16-node example is shown in Fig. 4. Note that the path over-provisioning allows 4 injection and 2 ejection switches to be reached by each gateway, as opposed to 2 injection and 1 ejection switches without over-provisioning [17]. Further, each injection switch is shared by two gateways. This decreases the probability that a gateway will be denied transmission due to traffic on the network.

Since the injection and ejection switches require less functionality than the routing switches, their designs may be simplified (i.e., certain waveguides and microrings may be eliminated), resulting in reduced optical losses and lower switch footprints (Figs. 3b and 3c, respectively). Then, the losses encountered in each photonic component may be accrued across the longest path of the network (gateway switch to gateway switch). We assume 0.05 dB and 0.5 dB of loss for every waveguide crossing and microring-utilizing turn, respectively. The routing switch, for example, has an average loss of 0.5 dB, with a maximum loss of 0.7 dB, depending on the specific input and output ports utilized. The average loss through the injection switch is determined to be 0.36 dB, with a maximum loss of 0.55 dB (in the case of an injection into the network). Lastly, the average loss through the ejection switch is 0.25 dB, with a maximum of 0.55 dB. Additionally, the longest path through the network consists of 11 routing switches (10 of which require the maximum loss value), 3 injection and ejection switches each (injected/ejected once, and encountered twice each within the network), 23 uni-directional link crossings, and 36 bidirectional link crossings, which totals to only 13.7 dB of optical loss.

This falls well within the optical loss budget offered by current optical transceivers used in high-density interconnects, especially those which utilize off-chip laser sources, as the distributed-feedback (DFB) lasers can provide output powers more than 20 dB above the sensitivities of current silicon optical receivers. This signifies the practicality of the hybrid approach. These assumed loss values reflect improvements in currently reported devices, for which the insertion losses are not fundamentally limited [19]; the loss values will most likely be obtainable soon due to the

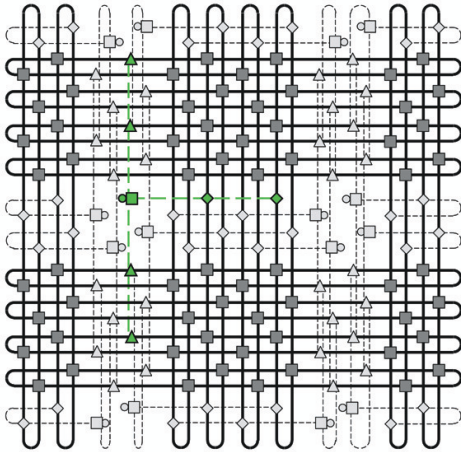


Figure 4. Layout of a 16-node 2-D folded torus with $2\times$ path over-provisioning. Solid (dashed) lines represent bi-directional (uni-directional) links. Network switches (squares) are shown in dark gray, and the GAPs, made of gateways (circles), gateway switches (squares), injection switches (triangles), and ejection switches (diamonds), appear in light gray. The GAP for one node is emphasized.

rapid advancement of photonic integration. Further, given the small footprint of the injection ($315\mu m \times 315\mu m$), ejection ($260\mu m \times 210\mu m$), and routing switches ($525\mu m \times 525\mu m$), the total area consumed by the photonic switches can be less than $1.4mm^2$ per node using a 1.6-nm wavelength grid; in addition, more area can be gained by using larger wavelength spacing.

Routing and Flow Control. Dimension order routing is a simple routing algorithm for mesh and torus networks, requiring minimal logic in the routers. We use XY dimension order routing on the torus network, with a slight modification required to accommodate the injection/ejection rules [17]: a source-encoded nested addressing scheme is employed to enforce the injection/ejection rules and exploit the path multiplicity in the torus network. Each message is encoded with two intermediate addresses guiding it through the GAPs and directing the routers in its path to route it to a selected path among the ones available.

The flow control technique in the network greatly differs from common NoC flow control methods. The dissimilarity stems from the fundamental differences between electronic and photonic technologies and mainly from the fact that memory elements (registers, SRAM, etc.) cannot be used to buffer messages or even to delay them while process-

ing is done. Electronic control packets are thus exchanged to acquire photonic paths, and the data are only transmitted, with a very high bandwidth, once the path has been acquired. Control packets are also used to tear down photonic paths after the message transmission is complete. The electronic control network can also be used for the exchange of short messages. Control that carry no data and are of a very small size such as read requests or cache snoops can be exchanged on the control network which is, in essence, a low-bandwidth electronic NoC. These control messages are not expected to create congestion in the control network because of their small size and do not require large resources in terms of additional circuitry or power. Further, some applications, such as cryptanalysis, for example, are characterized with the exchange of very small data messages without any locality that can be exploited for grouping or speculative fetching. A CMP featuring the proposed hybrid architecture can utilize the electronic control network to exchange these messages at a reasonable performance.

Network Interfaces. Electronic/Optical and Optical/Electronic (E/O and O/E) conversions are necessary for the exchange of photonic messages on the network. Each node includes a photonic network interface: a *gateway*. Small footprint microring resonator-based silicon optical modulators with data rates up to 12.5 Gb/s [20] as well as SiGe photodetectors [7] have been reported and have recently become commercially available [6], to be used in photonic chip-to-chip interconnects. The laser sources can be located off chip, externally coupled, as is typically the case in off-chip optical communication systems [6]. The network gateways also include the circuitry necessary for clock synchronization and recovery and serialization/deserialization.

Since electronic signals are fundamentally limited in their bandwidth to a few GHz, larger data capacity is typically provided by increasing the number of parallel wires. The optical equivalent of this wire parallelism can be provided by a large number of simultaneously modulated wavelengths using wavelength division multiplexing (WDM) [3] at the gateways. A recently reported translating device, implemented using microring resonator modulators, converts directly between space-parallel electronics and wavelength-parallel photonics in a manner that conserves chip space as the translator scales to very large data capacities [13].

The energy dissipated in these large parallel structures is not small, but it is still smaller than the energy consumed by the wide busses and buffers currently used in NoCs: the network gateway interface and corresponding E/O and O/E conversions occur once per node in the proposed system, compared to multiple ports at each router in electronic equivalent NoCs [16].

3 Path-Setup Procedure

The path acquisition procedure requires the path-setup packet to travel a number of electronic routers and undergo some processing in each hop. Contention may cause the packet to be blocked, leading to a path-setup latency on the order of tens of nanoseconds. Once a path is acquired, the transmission latency of the optical data is very short, depending only on the group velocity of light in a silicon waveguide: approximately 6.6×10^7 m/s, or 300 ps for a 2-cm path crossing a chip [9]. This latency mismatch is fundamental to intrachip optical communications: the network control and arbitration latency, determined by the electronic propagation velocity and processing speed, can impede the full exploitation of the latency advantages of optical transmission. This is independent of whether the control is performed in a centralized or distributed fashion. The path-setup procedure is, therefore, a key issue in determining the performance of the photonic NoC. Reductions in path-setup latency will directly translate to improved efficiency of the network interfaces, to higher average bandwidth, and to better exploitation of the optical medium.

For a given source-destination pair, the setup latency can be expressed as $D = (H-1) \cdot t_p + t_q$, where H is the number of hops in the packet's path, t_p is the processing latency in each router and t_q is the total additional latency due to contentions. Contentions in the path-setup phase are handled by queueing the path-setup packet until the message blocking its path is torn down and the path is cleared [17]. Simulations show that t_q is a major contributor to the overall setup latency, especially when the network is heavily loaded [17].

To reduce the contention-based setup latency, t_q , a new method of handling congestion is suggested. The new method relies on the fact that the actual processing latency in the path-setup phase, $(H-1) \cdot t_p$, is typically much lower than the contention-based latency. In the suggested methods, the buffering depth in the electronic router is reduced to zero. This means that when a path-setup packet is blocked, it is immediately dropped, and a *packet-dropped* packet is sent, on the control network, in the opposite direction to notify the sender. The sender can immediately attempt to set up an alternative path, exploiting the network's path multiplicity. With an adequate level of path-multiplicity, it is reasonable to assume that an alternative path can be found faster than it would take for the message obstructing the original path to be torn down.

Using the OMNET++ based POINTS simulator [17], a 36-core system with a photonic NoC and a path-multiplicity factor of $\times 2$ is simulated. The latency components in POINTS are based on predicted individual latencies of electronic and silicon-photonic components in a future 22nm process, and the optical message size is 16 KByte. The simulation results (Fig. 5) show that by setting the buffer depth




Figure 5. Average path-setup latency (top) and bandwidth (bottom) as a function of buffer depth in a 6×6 photonic NoC.

to 0, i.e. dropping every blocked packet and immediately notifying the sender, the path setup latency can be reduced by as much as 30%, when compared to the case where path-setup packets are not dropped on contention (buffer depth of 2). When a buffer depth of 1 is simulated, i.e. a single path-setup packet may be queued in each direction in each electronic router, the latency reduction is smaller.

The peak optical bandwidth per port in the simulations, using WDM and OTDM, is set at 960 Gb/s. The average bandwidth is calculated as the product of the peak bandwidth and the fractional time, in steady state, that can be allocated for actual transmission of the optical messages, after messages have been set up. The average bandwidth results are also shown in Fig. 5.

4 DMA Block Sizing

The Direct Memory Access (DMA) communication model is used in several interconnection network architectures requiring intensive bandwidth-driven interactions between processors. The Quadrics QsNetII [1] and the IBM Cell Element Interconnect Bus (EIB) [11] are two examples. DMA is appropriate for this application because it can be configured to use transactions of a fairly large, fixed size. This kind of bandwidth-intensive, large-transaction communication model is especially suitable for the photonic NoC design because it has the potential to exploit the




Figure 6. Average latency (top) and bandwidth (bottom) for network transaction of different sizes in a 6×6 photonic NoC.

large network bandwidth by reducing the fractional overhead of the path-setup procedure. Furthermore, the DMA overhead (requests, etc.) can be transmitted over the control network with a low latency while the optical path is being set up and thus some of the path-setup latency can overlap with the DMA overhead reducing total latency.

Accurate modeling of the DMA transaction requires the knowledge of the specific implementation of the DMA hardware [11]. However, the effects of the block size on the latency and on the average bandwidth in the network are simulated for the 6×6 network, for an unloaded network as well as for a heavily loaded network (offered load = 0.85). The peak transmission bandwidth is still 960 Gb/s. The results are shown in Fig. 6.

When inspecting Fig. 6 it is obvious that for small block sizes ($\leq 1KB$), the latency is dominated by the path-setup zero-load latency which is greater than the serialization latency, because of the extremely large transmission bandwidth. DMA blocks of this size will clearly be inefficient. When exceptionally large blocks are used ($\geq 65KB$) the increased serialization and contention latencies overshadow the gain in bandwidth which is diminishing for large blocks. The optimal block size for the transactions over the photonic NoC ranges, therefore, between 4 and 16 KB.

5 Conclusions

As multicore processors step into an era where high bandwidth communications is a key for computing performance, photonic NoCs offer a promising low-power solution, while presenting a challenging design problem. The suggested hybrid architecture of a photonic transmission NoC combined with an electronic control NoC addresses the design problems by using each technology according to its advantages: electronics for processing and photonics for transmission.

The physical layer arrangement of the photonic components is an important consideration for the performance of the network. Implementing the non-blocking photonic switches can increase the utilization, and simplifying the layout of the injection and ejection switches proves that the optical losses and area requirements of the network are realistic.

The path setup latency, identified as the key to high performance in the photonic NoC, is addressed in this paper and a technique that reduces it by 30% is presented and evaluated in simulation. As in other high-performance networks, DMA is adopted as an appropriate communication model because of the large block size and the possible overlap between the network overhead and the DMA overhead. Block sizes of 4KB and 16KB are identified as optimal for the network simulated in this study.

Acknowledgments

The authors would like to acknowledge the support of the U.S. Department of Defense under sub-contract B-12-664, and the National Science Foundation under grant CCF-0523771.

References

- [1] J. Beecroft et al. QsNet^{II}: defining high-performance network design. *IEEE Micro*, 25(4):34–47, July/Aug. 2005.
- [2] L. Benini and G. D. Micheli. Networks on chip: A new SoC paradigm. *IEEE Computer*, 49(2/3):70–71, Jan. 2002.
- [3] X. Chen, B. G. Lee, X. Liu, B. A. Small, I.-W. Hsieh, K. Bergman, J. Richard M. Osgood, and Y. A. Vlasov. Demonstration of 300 Gbps error-free transmission of WDM data stream in silicon nanowires. In *Conference on Lasers and Electro-Optics (CLEO 2007)*, May 2007.
- [4] W. J. Dally and B. Towles. Route packets, not wires: On-chip interconnection networks. In *Design Automation Conf.*, pages 684–689, June 2001.
- [5] W. J. Dally and B. Towles. *Principles and Practices of Interconnection Networks*. Morgan Kaufmann, San Francisco, CA, 2004.
- [6] C. Gunn. CMOS photonics for high-speed interconnects. *IEEE Micro*, 26(2):58–66, Mar./Apr. 2006.

- [7] A. Gupta, S. P. Levitan, L. Selavo, and D. M. Chiarulli. High-speed optoelectronics receivers in SiGe. In *17th Intl. Conf. on VLSI Design*, pages 957–960, Jan. 2004.
- [8] S. Heo and K. Asanović. Replacing global wires with an on-chip network: a power analysis. In *Intl. Symp. on Low Power Elect. and Design (ISLPED 2005)*, pages 369–374, Aug. 2005.
- [9] I.-W. Hsieh, X. Chen, J. I. Dadap, N. C. Panoiu, J. Richard M. Osgood, S. J. McNab, and Y. A. Vlasov. Ultrafast-pulse self-phase modulation and third-order dispersion in si photonic wire-waveguides. *Optics Express*, 14(25):12380–12387, Dec. 2006.
- [10] J. Kim, W. J. Dally, B. Towles, and A. K. Gupta. Microarchitecture of a high-radix router. In *ISCA '05: 32nd annual international symposium on Computer architecture*, pages 420–431, June 2005.
- [11] M. Kistler, M. Perrone, and F. Petrini. Cell multiprocessor communication network: Built for speed. *IEEE Micro*, 26(3):10–23, May/June 2006.
- [12] R. Kumar, V. Zyuban, and D. M. Tullsen. Interconnections in multi-core architectures: Understanding mechanism, overheads, scaling. In *ISCA '05: 32nd annual international symposium on Computer architecture*, June 2005.
- [13] B. G. Lee, B. A. Small, Q. Xu, M. Lipson, and K. Bergman. Characterization of a 4×4 Gb/s parallel electronic bus to WDM optical link silicon photonic translator. *IEEE Photon. Technol. Lett.*, 19(7):456–458, Apr. 1 2007.
- [14] T. M. Pinkston and J. Shin. Trends toward on-chip networked microsystems. *Intl. J. High Performance Computing and Networking*, 3(1):3–18, 2001.
- [15] R. Ramaswami and K. N. Sivarajan. *Optical Networks: A Practical Perspective*. Morgan Kaufmann, San Francisco, CA, second edition, 2002.
- [16] A. Shacham, K. Bergman, and L. P. Carloni. Maximizing GFLOPS-per-Watt: High-bandwidth, low power photonic on-chip networks. In *P=ac² Conference*, pages 12–21, Oct. 2006.
- [17] A. Shacham, K. Bergman, and L. P. Carloni. On the design of a photonic network on chip. In *The 1st IEEE International Symposium on Networks-on-Chips (NOCS 2007)*, May 2007.
- [18] S. Vangal et al. An 80-tile 1.28 TFLOPS network-on-chip in 65 nm CMOS. In *International Solid State Circuits Conf.*, Feb. 2007.
- [19] F. Xia, L. Sekaric, and Y. A. Vlasov. Ultracompact optical buffers on a silicon chip. *Nature Photonics*, 1:65–71, Jan. 2007.
- [20] Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson. 12.5 Gbit/s carrier-injection-based silicon microring silicon modulators. *Optics Express*, 15(2):430–436, 22 Jan. 2007.