Journal of
Nanoscience and Nanotechnology
Vol. 10, 1616–1625, 2010

A Macrochip Interconnection Network Enabled by Silicon Nanophotonic Devices*

Xuezhe Zheng[1,*], John E. Cunningham[1], Pranay Koka[2], Herb Schwetman[2], Jon Lexau[3], Ron Ho[3], Ivan Shubin[1], Ashok V. Krishnamoorthy[1], Jin Yao[4], Attila Mekis[4], and Thierry Pinguet[4]

[1]*Sun Microsystems, CTO Physical Sciences Center, San Diego, CA 92121, USA*
[2]*SUN Laboratories, Austin, TX 78727, USA*
[3]*SUN Laboratories, Menlo Park, CA 94025, USA*
[4]*Luxtera Inc., 2320 Camino Vida Roble, Carlsbad, CA 92011, USA*

We present an advanced wavelength-division multiplexing point-to-point network enabled by silicon nanophotonic devices. This network offers strictly non-blocking all-to-all connectivity while maximizing bisection bandwidth, making it ideal for multi-core and multi-processor interconnections. We introduce one of the key components, the nanophotonic grating coupler, and discuss, for the first time, how this device can be useful for practical implementations of the wavelength-division multiplexing network using optical proximity communications. Finite difference time-domain simulation of the nanophotonic grating coupler device indicates that it can be made compact (20 μm × 50 μm), low loss (3.8 dB), and broadband (100 nm). These couplers require subwavelength material modulation at the nanoscale to achieve the desired functionality. We show that optical proximity communication provides unmatched optical I/O bandwidth density to electrical chips, which enables the application of wavelength-division multiplexing point-to-point network in macrochip with unprecedented bandwidth-density. The envisioned physical implementation is discussed. The benefits of such an interconnect network include a 5–6×, improvement in latency when compared to a purely electronic implementation. Performance analysis shows that the wavelength-division multiplexing point-to-point network offers better overall performance over other optical network architectures.

Keywords: Optical Proximity Communications, Grating Coupler, Macrochip, Silicon Photonics, Nanophotonics, WDM Network, Optical Interconnects.

RESEARCH ARTICLE

1. INTRODUCTION

For reasons of increased performance, reduced power, and lowered design complexity, high performance processor architectures are increasingly integrating multiple processing cores on a single chip.[1,2] These chips exploit parallelism through multi-core and multi-thread processing, but ultimately, their ability to continue to scale processor performance depends on the ability of their on-chip interconnects networks to keep pace. Computers, even at the processor level, are now performance-limited by their networks. As designers continue to scale up the number of threads and cores integrated on a single chip, they must face conflicting demands. On one hand, more cores require higher bandwidth, and on the other hand higher integration limits the available silicon real estate and dissipated power. Future chip architectures are hence challenged to provide a low-latency, low-power, and high bandwidth network that can effectively interconnect the cores, and can keep pace with the computational performance improvements from technology scaling.

Even more challenging, the ever-increasing demands generated by new applications from military and commercial sectors require successively more computational power integrated on chip at a faster pace. Unfortunately, yield and lithographic reticle limits eventually constrain the number of processing cores on a single silicon chip. There is a need to break through and scale beyond the single reticle limit so that designers can extend their interconnect network with a design that can intimately integrate multiple chips together. This, however, would require much higher bandwidth-density I/Os for these chips to

*Author to whom correspondence should be addressed.
*This is an invited paper.

1533-4880/2010/10/1616/010 doi:10.1166/jnn.2010.2042

enable a network to interconnect them coherently with bandwidth matching the enormous aggregated processing power.

Scaling the required off-chip communication bandwidth by traditional electrical signaling methods such as SerDes will consume significant processor real estate and power. Electrical Proximity Communication, PxC, using capacitive coupling, is a more promising approach to alleviate part of the bottleneck with its combination of high bandwidth density and low energy per bit.[3,4] Moreover, the latency of an electrical path across multiple chips connected through PxC is set by repeated on-chip wires and is thus typically limited to velocities under 1/10 the speed of light,[5] which ultimately limits system performance. Optical signaling based on silicon photonics, on the other hand, could potentially eliminate the off-chip bandwidth bottleneck as well as provide lower latency communication than on-chip wires. Recent advances in silicon nanophotonics offer a potentially disruptive networking solution with unique advantages in low latency and high bandwidth while maintaining minimal power dissipation.[6–11] CMOS-compatible high speed, low power compact modulators,[12–16] detectors[15–18] and low-loss silicon waveguides enable electrical circuits to communicate optically through co-integrated photonic I/O interfaces and waveguide routing without the need of repeating, regeneration, or buffering.[19,20] With wavelength-division multiplexing (WDM) optics[21,22] they also offer unmatched bandwidth density. Furthermore, optical proximity communications (OPxC)[23,24] enable a seamless transparent I/O interface to bridge on-chip and off-chip bandwidth, and to naturally extend the logical boundary of a silicon chip beyond its economical reticle limit; in addition, they also enable multi-layer transparent optical routing. Taking advantage of these silicon nanophotonic technology advancements, designers can implement high performance optical networks for on-chip interconnections that span well beyond a chip's physical boundary for multi-chip communications.

In this paper, we will present an optical communication network that enables such evolutionary integration: a WDM point-to-point network enabled by novel silicon nanophotonic devices. This network offers strictly non-blocking all-to-all connectivity while maximizing bisection bandwidth, making it ideal for multi-core and multi-processor interconnections. In the following section, we will discuss the key nanophotonic device that enables practical implementation of the WDM network: the grating coupler-based OPxC. We will further show that such a photonic network can be extended to a "macrochip"[25] application naturally with very promising performance. The envisioned physical implementation will be discussed, and preliminary performance evaluation results will be reported. A summary section will follow.

2. SILICON PHOTONIC WDM POINT-TO-POINT NETWORK FOR MULTI-PROCESSOR INTERCONNECTS

For multiprocessor applications, the interconnect network must provide low power, low latency, high bandwidth and high-density communication between cores. The breakthrough in co-integrating high performance photonics with electrical circuits in a CMOS-compatible Si process[10] makes it practical and beneficial for optical signaling to replace on-chip global wiring for multi-core processor interconnects. Several networking architectures have been reported that exploit the large aggregate bandwidth and the high density of WDM optical interconnects available with the latest active and passive silicon photonic components. An electronically-reconfigurable circuit-switched mesh or torus architecture was introduced in Ref. [26], while an all-optical arbitration system based on ring modulators and add/drop filters was proposed in Ref. [27]. Both require system provisioning to establish a link path between the source and destination, and hence incur some extra communication latency. In addition, optical signals passing through many ring resonator based switches or add/drop filters may suffer from channel narrowing effects of cascaded WDM filters and crosstalk induced by waveguide crossings. The photonic network we introduce here employs a point-to-point topology, establishing all-to-all static links using WDM wavelength routing with the fewest components. It avoids crosstalk by employing a Manhattan routing topology using two separate optical layers. The routing topology for a sample 4×4 network with is shown in Figure 1.

Each site in the static WDM point-to-point optical network uses a combination of wavelength-division multiplexing in a waveguide and space-division multiplexing across multiple waveguides to establish a unique link to every other site in the macrochip. These are shown as 16 waveguides, each carrying one wavelength and each destined for a different target site. We draw all sixteen for illustrative reasons only; in a real system, we would only need fifteen optical transmitters to reach all other sites. A WDM multiplexer (Mux) merges groups of four transmitter (Tx) waveguides into a single four-wavelength WDM waveguide. This runs horizontally over to the column containing the destinations of those four Tx waveguides. Through a coupler, this 4-wavelength bundle coupled in the waveguide runs vertically. At each of the four destination sites, a WDM drop filter selects the appropriate wavelength, and drops it to the target receiver. For more bandwidth or accommodating more sites, one can simply use more wavelength channels and more waveguides. The network has no in-line switching and is non-blocking, yielding much lower latency and higher sustained bandwidth than hybrid optical networks employing dynamic, electronic switch configuration.

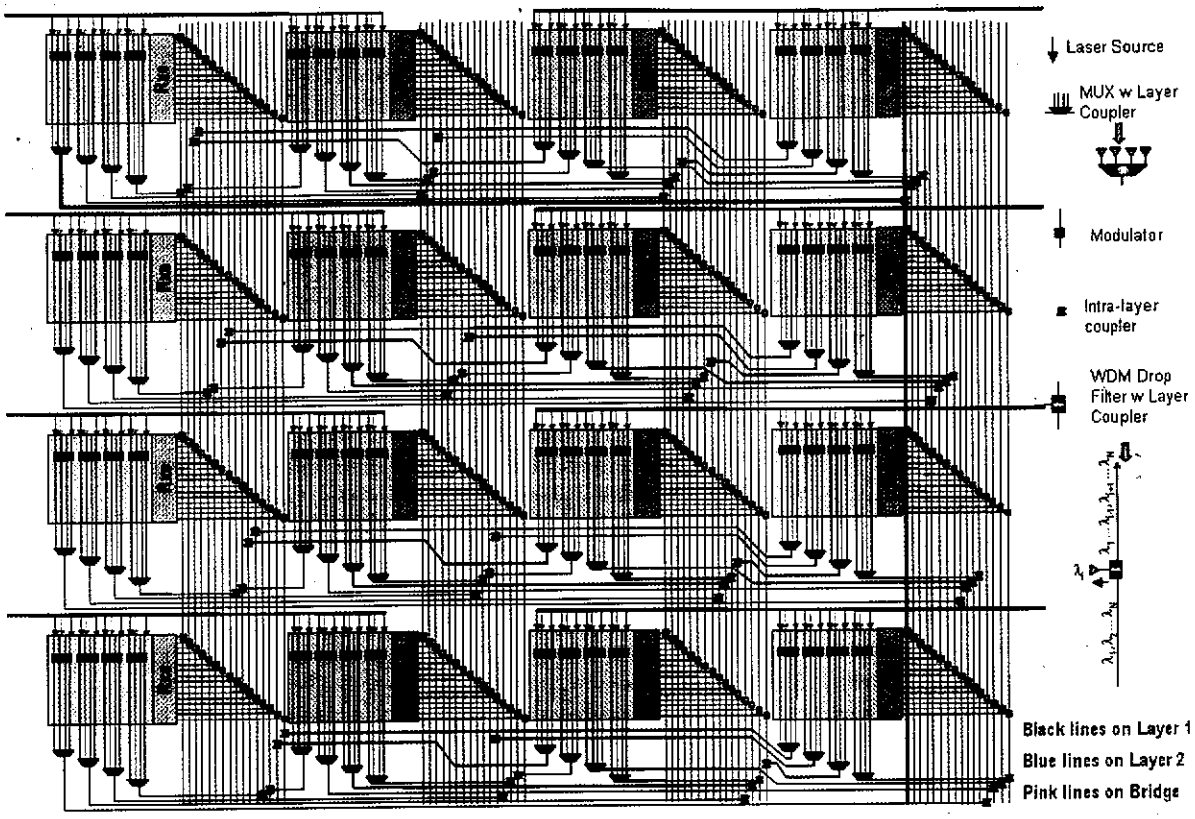

Fig. 1. A 4 × 4 WDM point-to-point network with base line connectivity.

This seemingly complicated network is actually made of fairly simple building blocks as shown in Figure 2. At the source site, a group of wavelength channels are multiplexed into one "row" waveguide, which is then routed to its destination column, and coupled into a column waveguide. The wavelength channels in the column waveguide are then routed and dropped to their destination sites respectively. The optical links are quite simple with a minimum number of photonic components involved for a complete WDM connection: modulator; waveguides; WDM multiplexer; waveguide couplers; distributed WDM de-multiplexers; and optical receiver. Optical amplifiers are typically not affordable for link power sensitive applications like multi-processor interconnects, and are not considered because such elements are not currently efficient. This essentially translates to a limited power budget for the optical links within the optical network. Optical links with fewer optical components are thus advantageous for their improved feasibility for practical implementations.

A few concerns arise when applying the WDM point-to-point network to large scale, high-bandwidth networks. These include waveguide crossings and the

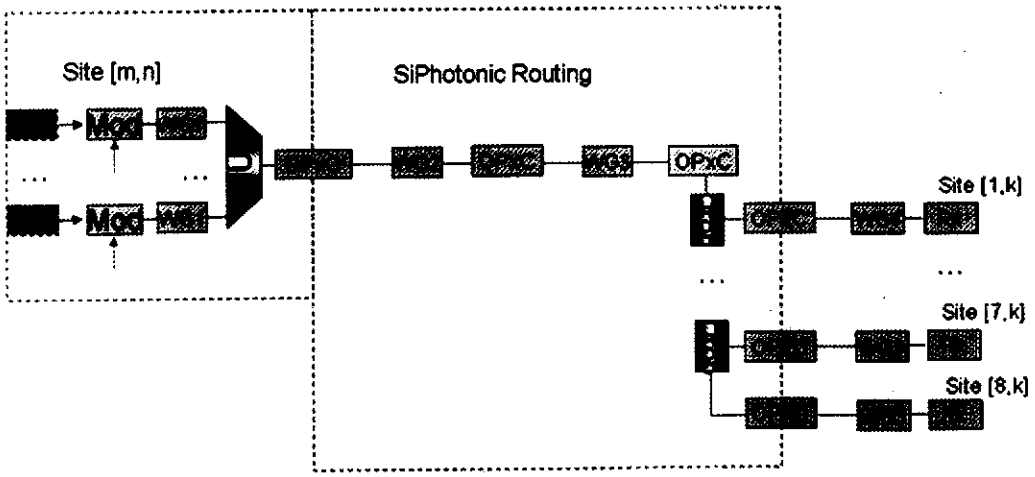

Fig. 2. Building block diagram of WDM point-to-point network.

area required for the photonic components themselves. Although nano-scale Si waveguide crossings can be implemented with relatively low loss and low crosstalk,[28,29] the aggregated effect could be fatal when potentially thousands of crossings inhabit the same chip. Similarly, even though any single Si waveguide has width less than 1 micron, a certain separation, typically up to 5 microns, may need to be maintained between waveguides to avoid signal crosstalk from neighboring waveguides. With thousands of waveguides for signal routing, significant routing area is needed for the network. One possible solution is to have a separate layer for photonics as suggested in Ref. [30], and 3D stacking of electronics with the photonic interconnect network. Another possibility is to use OPxC technology.[23,24] With OpxC, the photonic layer can further be separated into two layers to create more routing space and simultaneously to avoid waveguide crossings. One layer faces up to accommodate the photonic transmitters, receivers, wavelength multiplexers, and the East-West row waveguides , and a second layer can face down to house the North-South column waveguides (blue lines in Fig. 1) and wavelength demultiplexers. The two layers are interconnected via the OPxC couplers, as in the cartoon shown in Figure 3. The accurate alignment required for assembling the two layers together into one optical interconnect network can be achieved by registering sapphire balls in CMOS-compatible etched pits as demonstrated and reported in Ref. [31].

3. KEY ENABLING TECHNOLOGY: OPxC BASED ON A NANOPHOTONIC GRATING COUPLER (NGC)

As discussed above, the optical proximity communication coupling is a critical enabling component for practical implementation of the photonic network for multi-processor interconnects. OPxC couplers based on etch pits and reflecting mirror have been reported and showed promising performance.[23,24] However, to use it for an application with sub-micron waveguides, 3D waveguide tapers are required. It is therefore challenging to implement it with a fully CMOS-compatible process. Here, we introduce a new type of OPxC based on CMOS-compatible nanophotonic waveguide grating structures. Gratings have typically been used for coupling between waveguides and fibers.[32,33] Using two NGCs laid one on top of another in close proximity, OPxC coupling between two chips can be achieved, as conceptually shown in Figure 4.

Figure 5(a) depicts the structure of the grating coupler. Sub-micron Si waveguide is connected to the grating

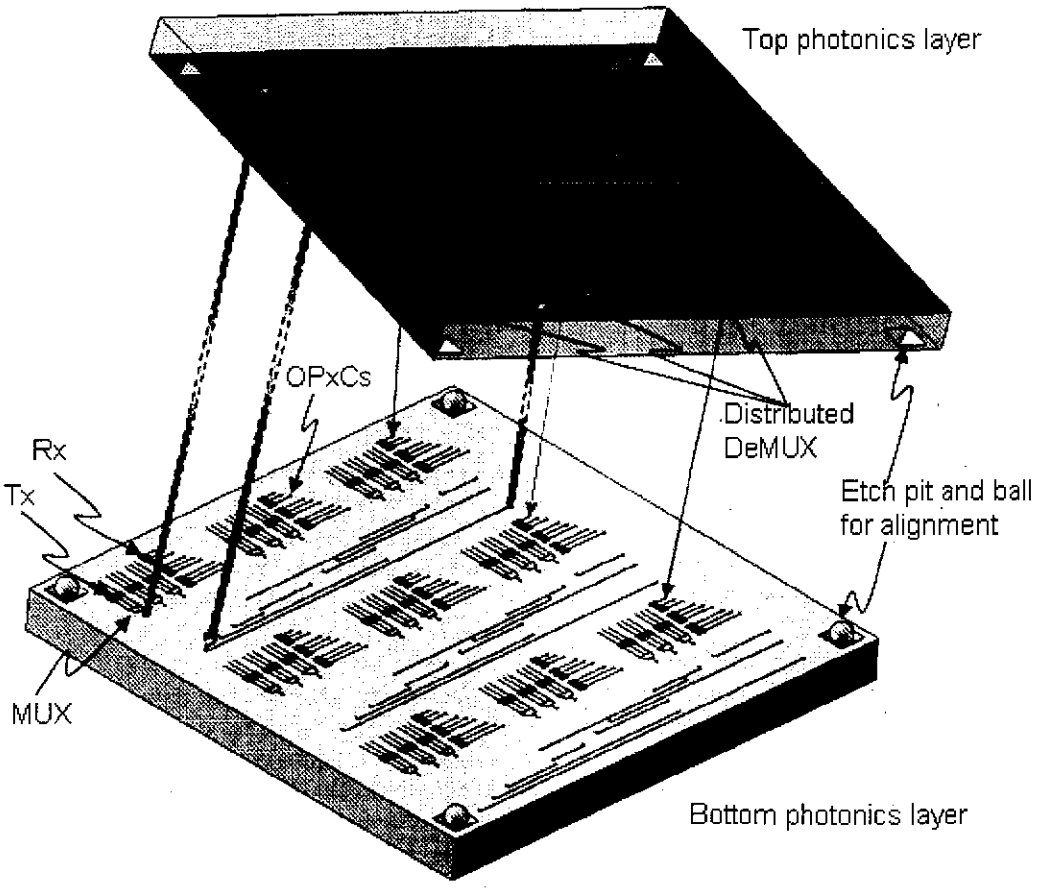

Fig. 3. Nanophotonics enabled WDM point-to-point network for on-chip multi-processor interconnects with double photonics layers.

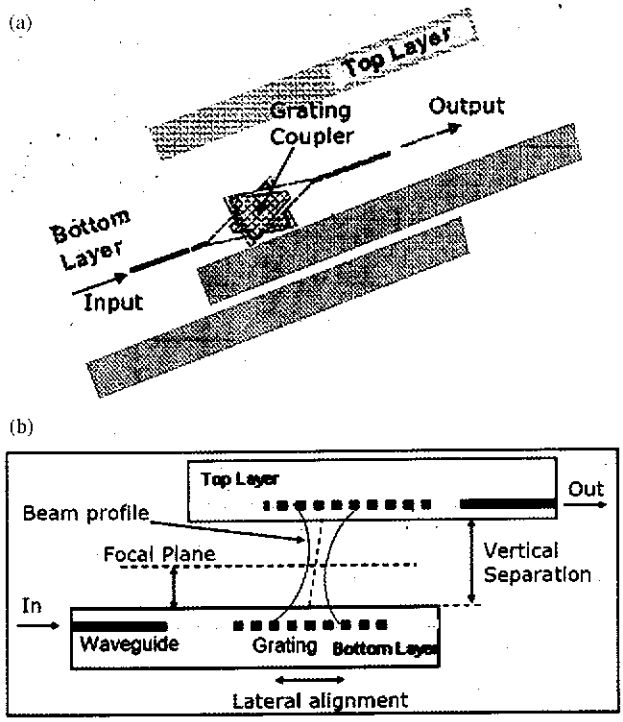

Fig. 4. Schematic of NGC to NGC OPxC structure.

structure through adiabatic taper. Since vertical coupling results in large second order reflections, an 8° from normal coupling angle is chosen for OPxC coupling, as shown in Figure 4(b). The nanophotonic grating structures are enabled with a beam focusing capability to generate a focused beam with its focal point halfway between the two chips. With the waveguides on the two chips terminated with the same grating couplers and aligned as illustrated in Figure 4(b), optical signals from the waveguides of one chip can be coupled to the waveguides on the other one across the inter-chip gap.

We used a finite-difference time-domain (FDTD) analysis to design the optimum grating structure. The SOI substrate thickness for the Si waveguides and grating couplers is optimized to provide high directionality towards the top surface of the chip. The buried oxide (BOX) thickness is selected appropriately to enhance reflection from the substrate in order to minimize optical loss from downward scattered light. To achieve optimum waveguide-to-waveguide coupling, the grating couplers are optimized to generate approximately symmetric mode profiles along both directions. A mode size of 10 μm is selected as the target at the focal plane because it is the size that can be achieved with the capabilities of 193 nm deep-UV lithography (~100 nm minimum feature size). The optimized GC has size of approximately 50 μm × 20 μm. Its mode profile at the focal plane along the x direction for OPxC coupling with 30 μm separation is shown in Figure 5(b). The mode profile is not perfectly Gaussian because the limited ability to produce small trenches in a real process

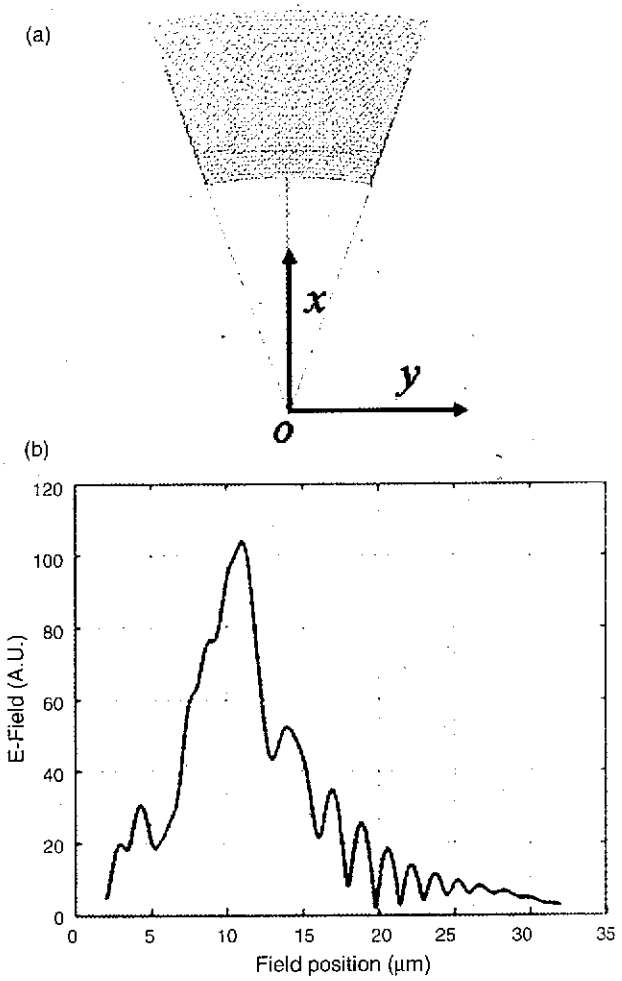

Fig. 5. Nanophotonic grating coupler. (a) Schematic view of grating coupler nanostructure. (b) Mode field profile of the NGC from the FDTD simulation at a separation of 30 microns along the x axis.

restricts the minimum scattering strength of the individual trenches. The resulting coupling loss of about 3.8 dB at a 10 μm separation includes the diffraction loss to the substrates when both grating OPxC couplers couple face-to-face. The loss is relatively high compared to GC to fiber coupling. It is because in the OPxC coupling configuration twice the substrate loss is incurred. In addition, there is more mode mismatch loss when two GCs are placed face-to-face in contrast to a GC coupling to a fiber because of the slightly asymmetric mode profile of the GC.

The grating structure is wavelength sensitive, and typically has angular dispersion that makes the output beams at different wavelengths propagate at different angles. For OPxC coupling with a finite separation, wavelength signals different from the designed center wavelength will have lateral position and angular offsets that cause higher coupling loss and limits the coupling bandwidth. Figure 6 shows the spectra of optimized GC OPxC coupling obtained from FDTD simulations for different separations. The coupling bandwidth depends on separation. With smaller separation, the lateral shift due to the angular

RESEARCH ARTICLE

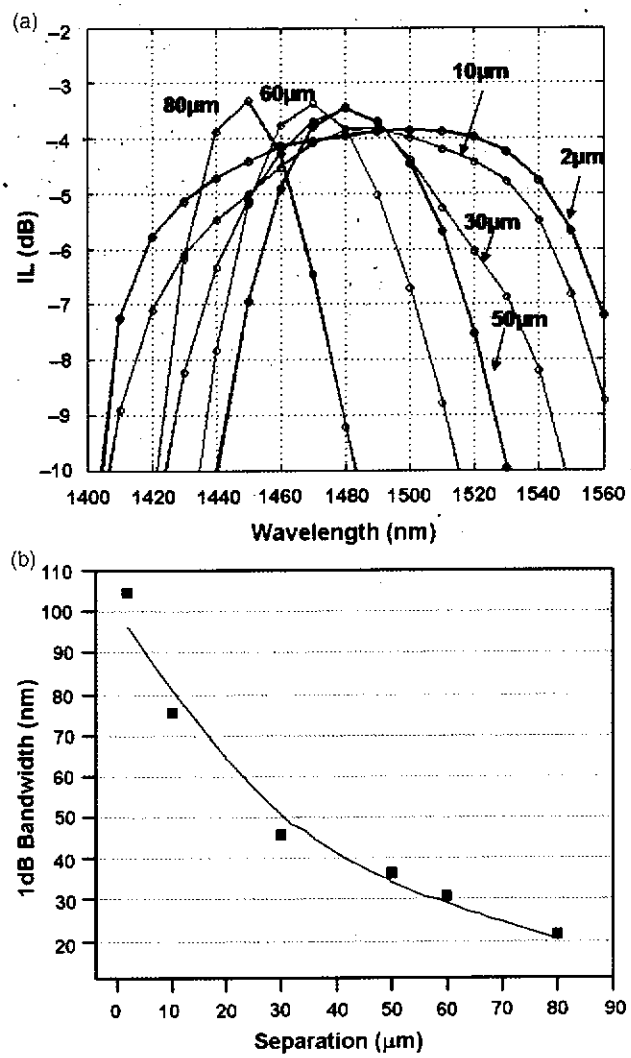

Fig. 6. Spectral bandwidth of a GC OPxC. (a) Spectra from FDTD simulation at optimized lateral positions for different separations. (b) 1 dB bandwidth for different chip separations.

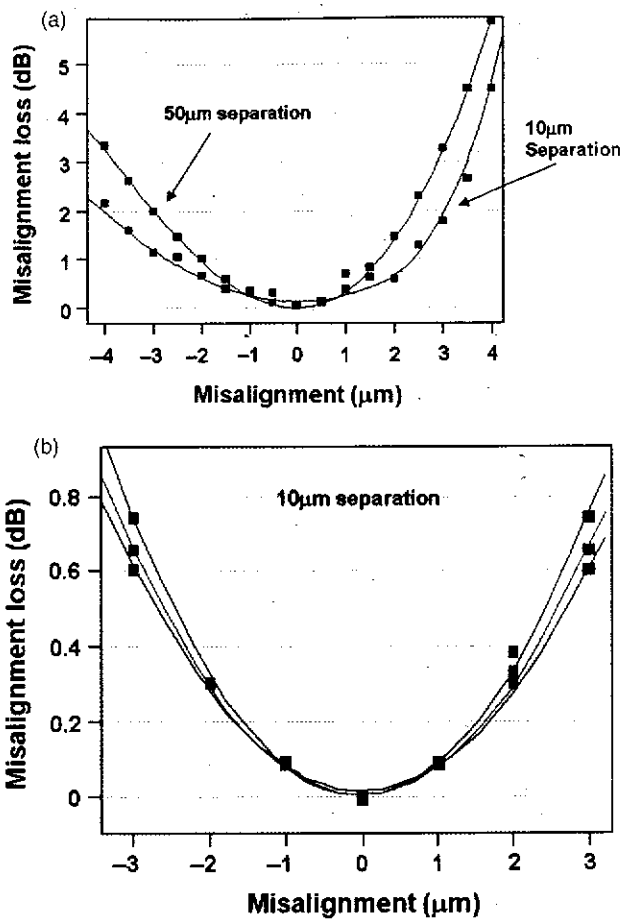

Fig. 7. Misalignment tolerance of GC based OPxC coupling. (a) Along x direction. Red line for 10 μm separation, and blue line for 50 μm separation. (b) Along y direction for 10 μm separation. Red, green and blue lines are for 1450 nm, 1490 nm and 1530 nm respectively.

dispersion of the grating is smaller, which therefore yields a larger coupling bandwidth. In contrast, at larger gaps, the coupling loss goes up quickly due to the larger walk-off for wavelength away from the center wavelength. The results also indicate that large coupling bandwidth, up to 100 nm, can be achieved with small chip separations of about 2 μm, as shown in Figure 6(b). The reason that Figure 6(a) shows different optimum center wavelengths for different gaps is because the optimization routine that simulated the alignment of the two chips was set to find the lowest insertion loss across the spectrum. At the same mode size, the beam divergence due to the diffraction is smaller for shorter wavelength. Therefore for larger gaps, the optimum operation wavelength (with the lowest possible loss) shifts towards shorter wavelengths due to their relatively longer collimation distances. If the realignment is instead done to optimize the positions of the NGCs according to the gap and the designed working angle, the optimum wavelengths for various gaps would have been the same as the original designed center wavelength, which in this case was set at 1490 nm.

Figure 7(a) shows the misalignment loss along the x direction, i.e., the symmetry axis of the NGC. At a 10 μm separation, a 1 dB excess loss is reached with misalignment of 2.5 μm. It is noted that the misalignment loss is not exactly symmetric. This is expected for the NGC OPxC coupling because of the non-ideal Gaussian beam shape formed by the GCs. Figure 7(b) shows the misalignment loss along the y direction, i.e., perpendicular to the symmetric axis of the NGC. The loss result is symmetric, as expected due to the symmetry of the NGC structure in this direction by design. At 10 μm separation, less than 0.8 dB excess loss is observed with misalignment of 3 μm. The small gap and accurate relative position tolerance required for large bandwidth, low loss NGC OPxC coupling is well within the capability of reported CMOS-compatible ball and etch pit alignment technology.[31]

Due to its small size and vertical coupling nature, many NGC couplers can be integrated in small area, offering electrical chips with enormous optical I/O bandwidth

when combined with WDM optics. Conservatively assuming 10 Gbps per wavelength channel and a channel spacing 1.6 nm (approx. 200 GHz), a NGC OPxC with a 50 nm bandwidth could potentially provide I/O bandwidth density up to 75 Petabitsps/cm^2. The performance of a NGC OPxC can potentially be improved in several ways. The current design employs two identical NGCs. There is large mode mismatch when they couple face-to-face due to their asymmetric mode profile. To improve the mode matching, one can use a pair of conjugated NGCs instead, one with $+8°$ angle while the other with $-8°$ angle. The loss performance can be also improved by mechanisms that facilitate upward reflection of light towards the NGC. For instance, one could apply distributed feedback reflector (DBR) layers between the bottom oxide and the silicon handle. In addition, the current design uses feature size available from CMOS 90 nm node process. In a more advanced CMOS process node, the symmetry of the NGC mode and its resemblance to the Gaussian mode profile can be further improved by redesigning the grating. Such a redesign is enabled by a smaller minimum feature size, which permit weaker scattering strengths at the beginning of the grating section, and can therefore result in lower insertion loss. We expect better than 1.5 dB coupling loss when these improvements are implemented.

4. SCALING BEYOND THE SINGLE RETICLE LIMIT

Single chip performance scaling is seriously challenged by limits on design complexity, yield, power, and limited on-chip real estate. A "macrochip" design integrates many chips into a logically contiguous large silicon and offers unprecedented computational density and energy efficiency.[25] One of the most critical components for a multiple-die macrochip is a high performance network that provides high bi-section bandwidth, low power, and low latency interconnects. Enabled by the nanophotonic OPxC discussed above, the WDM point-to-point network architecture can be scaled beyond single chip communications to macrochip interconnects seamlessly. As shown in Figure 8, electrical chips containing processor cores and memories are equipped with Si photonics for optical signaling, and arranged into a 2D array. OPxC couplers enable these opto-electronic high bandwidth density optical I/Os. Similar to the single chip networking discussed above, two separate optical routing layers with waveguides are added to physically interconnect the source and destination chips. The bottom routing layer facilitates the routing waveguides along the row direction, while the top routing layer (strips) houses the routing waveguides going along the column direction, and also the distributed de-multiplexers for dropping different wavelength channels at their destination sites. With this two layer routing structure, waveguide crossings can be easily avoided. A sample signal path of the WDM point-to-point network for the upper left corner chip to communicate with the bottom right corner chip is shown in Figure 8 with the red-arrow lines. Essentially, a group of wavelength channels on the source chip are multiplexed into one waveguide, coupled optically to the bottom routing layer via OPxC coupling, and then routed by the waveguides to its destination column; a slightly different scenario is to have the WDM multiplexers on the bottom routing layer. Next, a particular

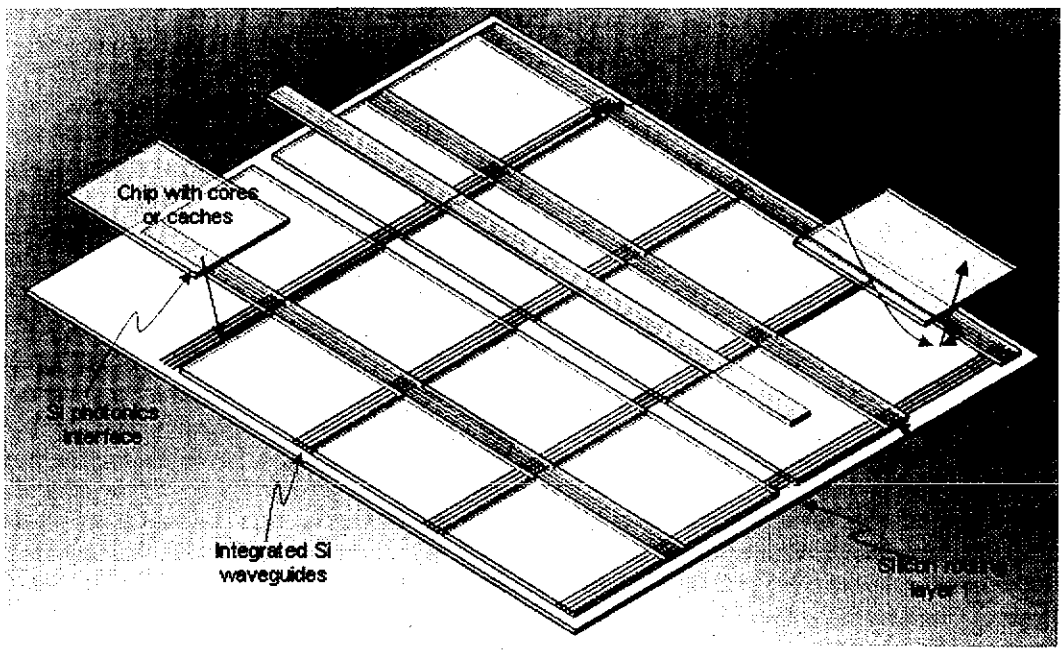

Fig. 8. Macrochip interconnected with Si photonic WDM point-to-point network.

RESEARCH ARTICLE

path is coupled to the top routing layer waveguide through OPxC, and further routed to its destination chip. A specific wavelength channel will then be dropped through a WDM drop filter. This particular channel will be coupled back to the bottom routing layer again through OPxC for transition, and finally received by the destination chip via its photonic interface. The remaining channels keep going forward until they reach their target sites. Similarly, WDM links can be added to interconnect his chip site to the chip sites in the other columns. Repeating the same process for all the chip sites in the array, the macrochip can be fully connected by this system-wide optical interconnect with low latency optical links.

5. PERFORMANCE ANALYSIS RESULTS

The WDM-based point-to-point network provides the maximum bisection bandwidth for a given number of transmitters and receivers at each site. The network bandwidth can be scaled up to physical limits by using more transmitter/receiver channels and WDM waveguides at each node, and the same routing approach can be applied to larger arrays as well. With each node equipped with 128×20 Gbps incoming and outgoing wavelength channels, grouped into 16 WDM waveguides, an 8×8 network can be constructed with a bisection bandwidth as high as 10 TB/s.

A static WDM point-to-point network provides dedicated communication channels between any pair of nodes at any given time, and enables a strictly non-blocking topology with no per-packet setup time. Figure 9 shows the basic latency characteristics of such a network, built with WDM SiP components for an 8×8 processor array, and compares them against a comparable electronic implementation using electrical proximity interconnect.[25] We observe that the WDM point-to-point network achieves superior performance over the electronic implementation

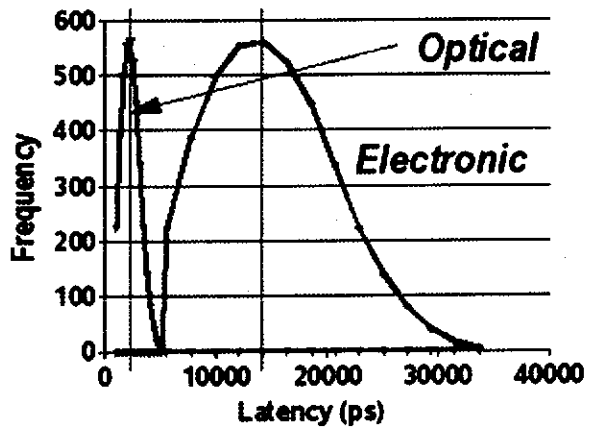

Fig. 9. Latency characteristic of WDM point-to-point network versus electrical implementation.

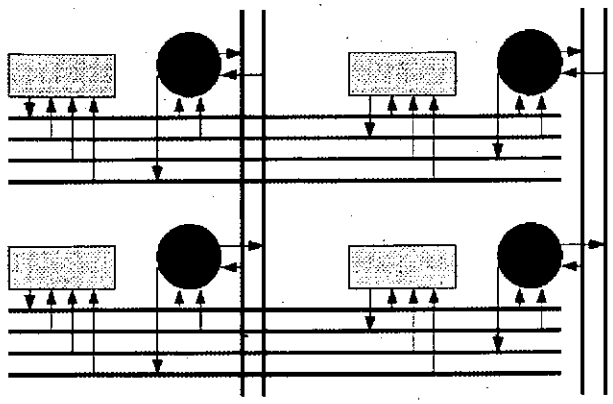

Fig. 10. A packet switching network with routers at each node.

with a 5–6× improvement in average latency and a much smaller variance in latencies across different paths.

We also performed a performance comparison of our non-blocking point-to-point interconnect against two alternative network architectures, a packet-switched network and a canonical hybrid optical mesh.

5.1. Packet Switched Network

Figure 10 shows the schematic of a 2×2 version of the packet switched network we simulated. The vertical and horizontal lines represent optical data channels. The routers are electronic routers requiring an optical–electrical conversion at the inputs and an electrical–optical conversion at the outputs. All nodes and routers have a broadcast channel along the row they belongs to. Apart from the row-broadcast channel, each router also has a broadcast channel along the column it belongs to. All communications between the nodes in the same row take place without the need of the router. Inter-row communications require two router hops.

5.2. Hybrid Optical Mesh

The hybrid mesh network provides a mesh of optical switches that are controlled by an overlying electronic router network. The hybrid mesh is non-blocking as long as no two source nodes communicate with the same destination node simultaneously. Each node is required to establish an optical circuit switched path to the destination node before it transmits data. The path setup and tear-down is performed using the electronic network. The hybrid mesh incurs a path setup latency, hence for transmission of small messages the overheads can be quite high.

5.3. Performance Methodology and Results

We developed an event driven simulation model of the three networks. We used a 320 GB/s of outgoing bandwidth per node and kept this constant across all network

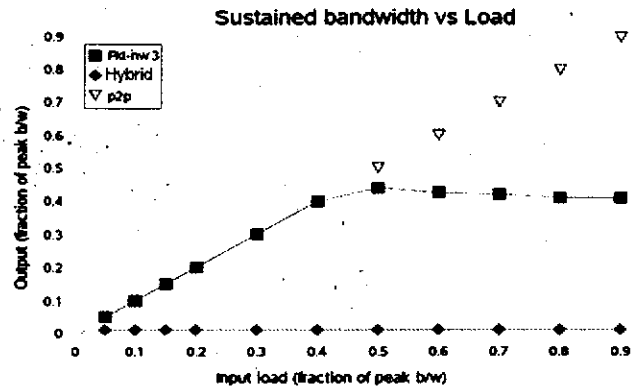

Fig. 11. Bandwidth response of point-to-point, hybrid mesh and packet switched networks.

simulations. We assumed a routing delay of 0.5 ns at the electronic routers in the packet switched networks and an optical–electrical conversion latency on 0.15 ns. We simulated an 8×8 array of nodes, with a total bandwidth of 20 TB/s. We show the performance results for 128 byte size messages. We pick a small message size because small messages are common to shared memory machines and several high-performance computing (HPC) benchmarks; also, performance at smaller message sizes is highly affected by the network overheads and clearly illustrates the difference between these network topologies. For all simulations we generated input load with a uniform probability to different destination nodes. Figure 11 shows the sustained bandwidth plot for all three networks.

The X-axis shows the input load or the bandwidths at which the network was driven, as a fraction of the peak bandwidth (20 TB/s). The Y-axis shows the output bandwidth response of the network as a fraction of the peak bandwidth. Our point-to-point network has no setup latencies and is completely non-blocking and hence the bandwidth response scales linearly with input load. The packet switched network on the other hand is limited by the bandwidth the router can handle and the processing latency in the electronic domain. Even with an aggressive serialization delay of 0.5 ns, the router can handle only up to 43% of the peak bandwidth. The hybrid mesh performs extremely poorly at small packets sizes. The electronic network and the routers are much slower than the optical counterparts and hence constitute a large path setup latency for every small packet. The hybrid mesh delivers only 1% of its peak bandwidth for small packet sizes. As can be seen from Figure 11, the WDM point-to-point network has much better bandwidth characteristics at random load than the other two networks.

6. SUMMARY

The interconnect network is becoming an increasingly important part of computing system design. In this paper, we presented a WDM point-to-point network based on Si nanophotonic devices for multi-processor interconnects. CMOS nano-photonics provide electrical circuits high bandwidth photonic I/O interfaces, enabling low power, low latency optical signaling for interconnects. OPxC based on nano-photonic grating couplers enables separate photonic routing layers to avoid waveguide crossings and to alleviate chip real estate limits. More importantly, it equips electrical chips with unmatched optical I/O bandwidth density, and consequently makes it possible to scale the WDM point-to-point interconnect network architecture beyond single reticle limits to achieve macrochip system interconnects. We showed that this non-blocking, point-to-point WDM routing network has superior performance and no setup delays when compared to both an electrically controlled switching network and a packet-switched network of the same bandwidth. This improvement is particularly evident as the loading of the network goes up. It further simplifies the control of the network and eliminates the resulting power required for network resource arbitration. The static WDM non-blocking network topology provides efficient transport for small messages (64 B or less), an important characteristic for supporting shared memory machines or HPC challenge benchmarks that make use of small messages.

Acknowledgments: This material is based upon work supported, in part, by DARPA under Agreement No. HR0011-08-09-0001. The authors would like to thank Dr. Brian O'Krafka, formerly of Sun Microsystems, for his early contribution to this work.

References and Notes

1. S. Manish, J. Barreh, J. Brooks, R. Golla, G. Grohoski, N. Gura, R. Hetherington, P. Jordan, M. Luttrell, C. Olson, S. Bikram, D. Sheahan, L. Spracklen, and A. Wynn, *Solid-State Circuits Conference, ASSCC '07 IEEE Asian*, IEEE, Jeju, Korea, November (**2007**), p. 22.
2. S. Vangal, J. Howard, G. Ruhl, S. Dighe, H. Wilson, J. Tschanz, D. Finan, P. Iyer, A. Singh, T. Jacob, S. Jain, S. Venkataraman, Y. Hoskote, and N. Borkar, *IEEE International Solid-State Circuits Conference 2007*, IEEE, ISSCC Digest of Technical Papers, San Francisco, February (**2007**), p. 98.
3. R. J. Drost, R. D. Hopkins, R. Ho, and I. E. Sutherland, *IEEE J. Solid-State Circuits* 39, 1529 (**2004**).
4. D. Hopkins, A. Chow, R. Bosnyak, B. Coates, J. Ebergen, S. Fairbanks, J. Gainsley, R. Ho, J. Lexau, F. Liu, T. Ono, J. Schauer, I. Sutherland, and R. Drost, *IEEE Int. Solid-State Circuits Conf.*, IEEE, ISSCC Digest of Technical Papers, San Francisco, February (**2007**), p. 368.
5. R. Ho, K. Mai, and M. Horowitz, *Proceedings of the IEEE* 89, 490 (**2001**).
6. R. Soref, *Proceedings of the IEEE* 72, 1687 (**1993**).
7. L. C. Kimerling, Towards the First Silicon Lasers, edited by L. Pavesi et al., Kluwer Academic Publishers, Boston (**2003**), p. 465.
8. B. Jalali and S. Fathpour, *IEEE/OSA Journal of Lightwave Technology* 24, 4600 (**2006**).

RESEARCH ARTICLE

9. R. Soref, *IEEE J. Selected Topics Quantum Electron.* 12, 1678 (**2006**).
10. C. Gunn, *IEEE Micro* 26, 58 (**2006**).
11. M. Lipson, *J. Lightwave Technol.* 23, 4222 (**2005**).
12. J. Liu, M. Beals, A. Pomerene, S. Bernardis, R. Sun, J. Cheng, L. C. Kimerling, and J. Michel, Group IV Photonics, *5th IEEE International Conference*, IEEE, Sorrento, Italy, September (**2008**), p. 10.
13. T. Pinguet, V. Sadagopan, A. Mekis, B. Analui, D. Kucharski, and S. Gloeckner, Group IV Photonics, *4th IEEE International Conference*, IEEE, Tokyo, Japan, September (**2007**), p. 186.
14. Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, and M. Lipson, *Opt. Express* 15, 430 (**2007**).
15. A. Narasimha and B. Analui, *Proc. Optical Fiber Communications (OFC/NFOEC) Conf.*, IEEE, San Diego, February (**2008**).
16. G. Masini, G. Capellini, J. Witzens, and C. Gunn, Group IV Photonics, *4th IEEE International Conference*, IEEE, Tokyo, Japan, September (**2007**), p. 28.
17. T. Yin, R. Cohen, M. M. Morse, G. Sarid, Y. Chetrit, D. Rubin, M. J. Paniccia, *Optical Fiber Communication/National Fiber Optic Engineers Conference, OFC/NFOEC 2008*, IEEE, San Diego, February (**2008**), p. 1.
18. S. J. Koester, C. L. Schow, L. Schares, G. Dehlinger, J. D. Schaub, F. E. Doany, and R. A. John, *IEEE Journal of Lightwave Technology* 25, 46 (**2008**).
19. M. A. Webster, R. M. Pafchek, G. Sukumaran, and T. L. Koch, Low-loss thin SOI waveguides and high-q ring resonators, *Optical Society of America Annual Meeting*, Tucson, October (**2005**).
20. D. K. Sparacin, S. J. Spector, and L. C. Kimerling, *IEEE J. Lightwave Technol.* 23, 2455 (**2005**).
21. B. E. Little, S. T. Chu, H. A. Haus, J. Foresi, and J. P. Laine, *IEEE J. Lightwave Technol.* 15, 998 (**1997**).
22. B. G. Lee, B. A. Small, Q. Xu, M. Lipson, and K. Bergman, *IEEE Photon. Tech. Lett.* 19, 456 (**2007**).
23. X. Zheng, J. E. Cunningham, I. Shubin, J. Simons, M. Asghari, D. Feng, H. Lei, D. Zheng, H. Liang, C. Kung, J. Luff, T. Sze, D. Cohen, and A. V. Krishnamoorthy, *Optics Express* 16, 15052 (**2008**).
24. A. V. Krishnamoorthy, J. E. Cunningham, X. Zheng I. Shubin, J. Simons, D. Feng, H. Liang, C. Kung, and M. Asghari, *IEEE Journal of Quantun Electronics* 45, 409 (**2009**).
25. A. V. Krishnamoorthy, R. Ho, X. Zheng, H. Schwetman, J. Lexau, P. Koka, G. Li, I. Shubin, and J. E. Cunningham, *Proceedings of the IEEE* 97, 1337 (**2009**).
26. M. Petracca, B. G. Lee, K. Bergman, and L. P. Carloni, *Proc. 16th IEEE Symposium on High Performance Interconnects (HOTI '08)*, IEEE, Stanford, August (**2008**), p. 31.
27. R. G. Beausoleil, J. Ahn, N. Binkert, A. Davis, D. Fattal, M. Fiorentino, N. P. Jouppi, M. McLaren, C. M. Santori, R. S. Schreiber, S. M. Spillane, D. Vantrease, Q. Xu, *Proc. 16th IEEE Symposium on High Performance Interconnects (HOTI '08)*, IEEE, Stanford, August (**2008**), p. 182.
28. W. Bogaerts, P. Dumon, D. V. Thourhout, and R. Baets, *Opt. Lett.* 32, 2801 (**2007**).
29. P. Sanchis, J. V. Galan, A. Brimont, A. Griol, J. Marti, M. A. Piqueras, J. M. Perdigues, Group IV Photonics, *4th IEEE International Conference*, September (**2007**), p. 1.
30. A. V. Krishnamoorthy, A. Lentine, K. Goossen, J. Walker, T. Woodward, J. Ford, G. Aplin, L. D'Asaro, S. Hui, B. Tseng, R. Leibenguth, D. Kossives, D. Dahringer, L. Chirovsky, and D. Miller, *IEEE Photonics Technology Letters* 7, 1288 (**1995**).
31. J. E. Cunningham, A. V. Krishnamoothy, H. Eberle, N. Gura, I. Shubin, D. Hopkins, R. Drost, J. Mitchell, R. Ho, A. Chow, J. Lexau, J. Schauer, J. Shi, W. Olesinski, and J. Ebergen, Active Demonstration of a Passively Self-Aligned, Multi-Chip Package Using Proximity Communication in a Switching Fabric, *42nd International Symposium on Microelectronics, IMAPS 2009*, San Jose, November (**2009**), to be published.
32. L. Vivien, G. Maire, G. Sattler, D. Marris-Morini, E. Cassan, S. Laval, A. Kazmierczak, D. Giannone, B. Sanchez, A. Griol, D. Hill, K. B. Gylfason, and H. Sohlstrom, Group IV Photonics, *4th IEEE International Conference*, IEEE, Tokyo, Japan, September (**2007**), p. 162.
33. G. Z. Masanovic, V. M. N. Passaro, and G. T. Reed, *IEE Proc.-Optoelectron.* 152, 41 (**2005**).

Received: 31 December 2009. Accepted: 31 March 2009.