

UC Davis

UC Davis Previously Published Works

Title

Design and evaluation of an arbitration-free passive optical crossbar for on-chip interconnection networks

Permalink

<https://escholarship.org/uc/item/3xm1g01w>

Journal

Applied Physics A: Materials Science & Processing, 95(4)

ISSN

1432-0630

Authors

Zhou, Linjie
Djordjevic, Stevan S.
Proietti, Roberto
[et al.](#)

Publication Date

2009-06-01

DOI

10.1007/s00339-009-5121-6

Peer reviewed

Design and evaluation of an arbitration-free passive optical crossbar for on-chip interconnection networks

Linjie Zhou · Stevan S. Djordjevic · Roberto Proietti ·
Dan Ding · S.J.B. Yoo · Rajeevan Amirtharajah ·
Venkatesh Akella

Received: 5 September 2008 / Accepted: 16 December 2008 / Published online: 20 February 2009
© The Author(s) 2009. This article is published with open access at Springerlink.com

Abstract With recent advances in silicon nanophotonics, optical crossbars based on CMOS-compatible microring resonators have emerged as viable on-chip optical interconnection networks to deliver high-bandwidth communication at low power dissipation with a small footprint. This paper describes the design, fabrication and evaluation of an arbitration-free passive crossbar based on a microring resonator matrix that can be used to route wavelength division multiplexing (WDM) signals across the chip. The salient feature of the proposed design is the ability to support multicasting and many-to-one communication efficiently (without arbitration), which makes it suitable for implementing cache coherency protocols and on-chip interconnect in future many-core processors.

PACS 42.82.Ds · 42.82.Et · 42.79.Hp

1 Introduction

Optical interconnection networks are superior to their electrical counterparts in terms of power consumption and immunity from electromagnetic interference [1]. Recent developments in silicon photonics have demonstrated various complimentary-metal-oxide-semiconductor (CMOS)-

compatible devices suitable for on-chip optical interconnections. They include high-speed and low-power electro-optic switches, modulators, and detectors [2–15] exploiting low-loss silicon photonic waveguides. In particular, silicon photonic microring modulators have achieved modulation bandwidth beyond 10 GHz [12], and it is expected that a data rate of 40 GHz is possible in the future at very low energy consumption per bit [16]. Intel recently reported 31 GHz germanium photodetectors that can be integrated on silicon waveguides [14]. Hybrid bonding also has been investigated to provide laser sources on a silicon platform [3, 9]. Reference [17] highlights the fundamental benefits of nanophotonic global interconnects when compared to their electronic counterparts with respect to power consumption. Reference [18] compares packet switched electronic and photonic networks and concludes that a hybrid electronic/photonic network can dramatically reduce power consumption. In [19], researchers from MIT demonstrated the benefits of on-chip networks with monolithic silicon photonics and demonstrated that optical links can have significant advantages in terms of energy efficiency compared to an electrical link. In [20], researchers demonstrate the potential benefits of optical interconnects in the CPU/DRAM interface in high performance multicore processors.

In this paper, we demonstrate a microring resonator matrix-based optical crossbar that can be used in an on-chip optical interconnect network. The optical paths between the sources and destinations are established based on the wavelengths that are used. The proposed crossbar allows wavelength division multiplexing (WDM) signals from a single source to be multicast to an arbitrary subset of the output ports. Also, signals from multiple sources can be multiplexed to a single destination through the proposed crossbar without arbitration (henceforth referred to as arbi-

L. Zhou (✉) · S.S. Djordjevic · D. Ding · S.J.B. Yoo ·
R. Amirtharajah · V. Akella
Department of Electrical & Computer Engineering, University of
California, Davis, CA 95616, USA
e-mail: ljzhou@ucdavis.edu
Fax: +1-530-7528428

R. Proietti
Scuola Superiore Sant'Anna, Pisa 56124, Italy

tration-free) for the destination, as the incoming signals are distributed over distinct wavelength channels.

The rest of the paper is organized as follows. We begin with an overview of the operation of a single microring resonator crossbar (basic building block) and show how it can be composed to form more complicated $N \times N$ crossbar networks. Next, we use the transfer matrix method to model the crossbar and predict how the critical parameters will affect its transmission performances to guide our design. The following section shows the fabrication processes and presents the preliminary experimental results. Finally, this paper concludes by discussing the potential applications of the proposed crossbar network in emerging manycore processor arena.

2 Design overview

The basic building block for the optical crossbar is an add-drop filter comprising a microring resonator coupled with a waveguide crossing [19]. Figure 1(a) shows the schematic of the basic building block. The optical signal excited at the input port propagates along the vertical waveguide. When the signal wavelength does not correspond to the resonance wavelength of the microring resonator, the optical signal continues to propagate without tunneling to the other orthogonal (horizontal) waveguide. When the signal wavelength is at the resonance wavelength, the optical signal couples into the microring resonator and tunnels out to the other waveguide. In this way, the optical signal can be routed to either the through or drop port, depending on its wavelength.

The waveguide crossing junction usually causes high scattering loss, lowering signal forward transmission and inducing crosstalk between the two output ports. The optical crossbar exploits multimode interference (MMI) couplers to mitigate this problem [20]. For the MMI coupler, the

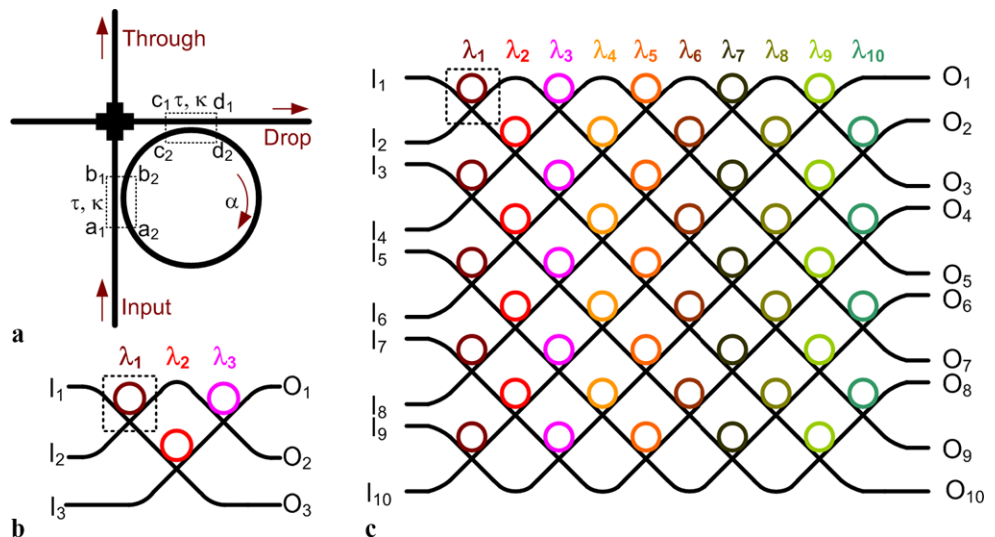
waveguide mode is self-imaged at the MMI center (crossing center) with no expansion of the wavefront, resulting in smoother forward propagation. Simulations indicate that the junction insertion loss can be as low as 0.1 dB per junction and crosstalk can be suppressed to -40 dB using a carefully designed MMI coupler [20].

Larger-scale crossbar designs require cascading multiple basic building blocks in a matrix form [21]. Figure 1(b) shows a 3×3 crossbar constructed using three basic building blocks. Figure 1(c) shows a 10×10 crossbar constructed using 45 basic building blocks. In order to establish an arbitration-free wavelength routing path between input and output ports, each column of the microring resonator matrix should have different resonance wavelengths. Thus, for an $N \times N$ crossbar, the microring resonators work at N different resonance wavelengths, and the total number of microrings is $N(N - 1)/2$. Table 1 illustrates the input-vs.-output routing table for a 3×3 crossbar with three microrings whose resonances are at λ_1, λ_2 and λ_3 . Here, three WDM signals (λ_1, λ_2 and λ_3) from one input-port are demultiplexed at the output after passing through the crossbar and the wavelength assignment is cyclic for different input ports. In this sense, the crossbar is functionally similar to arrayed waveguide gratings (AWG), except that it uses optical resonance rather than optical phase coherence to separate WDM channels. Compared with AWGs, the microring matrix-based crossbars are far more compact, given that microrings have diameters of a few microns [11]. Note that

Table 1 3×3 crossbar routing table

	O_1	O_2	O_3
I_1	λ_2	λ_1	λ_3
I_2	λ_3	λ_2	λ_1
I_3	λ_1	λ_3	λ_2

Fig. 1 (a) Add-drop filter basic building block, (b) 3×3 crossbar, and (c) 10×10 crossbar. a_m, b_m, c_m and d_m ($m = 1, 2$) are the electric field; τ and κ are the transmission and coupling coefficients of the couplers; α is the waveguide propagation loss in the ring



large-scale crossbar suffers more losses due to the cascaded waveguide crossing junctions.

3 Modeling

The modeling process is based on applying the transfer matrix method to the microring resonator-based crossbar [22]. The first step in the process is to deduce the electric field transmission functions for the two output ports of the basic building block. The electric field transmission through the input coupler of the microring resonator can be expressed as

$$\begin{bmatrix} b_1 \\ b_2 \end{bmatrix} = \begin{bmatrix} \tau & -i\kappa \\ -i\kappa & \tau \end{bmatrix} \begin{bmatrix} a_1 \\ a_2 \end{bmatrix}, \quad (1)$$

where a_m and b_m ($m = 1, 2$) are the electric field before and after the input coupler, τ the electric field transmission coefficient through the coupler, and κ is the electric field coupling coefficient. For lossless coupling, $\tau^2 + \kappa^2 = 1$. Assuming that the output coupler is identical to the input one, the electric field before and after the output coupler c_m and d_m are thus related by the same matrix in (1).

The electric field inside the microring resonator is related as

$$\begin{bmatrix} c_2 \\ d_2 \end{bmatrix} = \begin{bmatrix} A^{1/4}e^{-i\beta L/4} & 0 \\ 0 & A^{-3/4}e^{i\beta 3L/4} \end{bmatrix} \begin{bmatrix} b_2 \\ a_2 \end{bmatrix}, \quad (2)$$

where A is the fraction of the electric-field amplitude that remains upon a microring round trip, L is the microring resonator round-trip length, $\beta = n_{\text{eff}}k_0$ is the propagation constant in the microring resonator, n_{eff} is the microring waveguide effective index of refraction, and $k_0 = 2\pi/\lambda$ is the free-space propagation constant (λ is the free-space wavelength). A is related to the microring waveguide propagation loss α (in dB/cm): $\alpha = -8.68 \ln(A)/L$. The microring waveguide loss α consists of silicon material loss (including free carrier absorption loss) α_{si} , waveguide sidewall roughness induced scattering loss α_{sc} , and waveguide bending radiation loss α_{b} , i.e., $\alpha = \alpha_{\text{si}} + \alpha_{\text{sc}} + \alpha_{\text{b}}$.

Solving the above equations leads to the through and drop port electric field transmission functions, assuming waveguide crossing junction loss is negligible:

$$T_t \equiv \frac{b_1}{a_1} = \frac{\tau - A\tau e^{-i\beta L}}{1 - A\tau^2 e^{-i\beta L}}, \quad (3)$$

$$T_d \equiv \frac{d_1}{a_1} = \frac{-\kappa^2 \sqrt{A} e^{-i\beta L/2}}{1 - A\tau^2 e^{-i\beta L}}. \quad (4)$$

As the input and output couplers are identical, the microring resonator always works in the ‘under-coupling’ regime [22]. In order to transfer more power to the drop waveguide at resonance wavelengths, κ should be large and α should be small. Figure 2 shows the through transmission spectrum extinction ratio (ER) and resonance linewidth (3-dB bandwidth) contour maps as a function of κ and α . The extinction ratio at resonance is high when the coupling is close to the ‘critical-coupling’ point (small α and large κ), which favors the signal dropping to the other waveguide. The resonance 3-dB bandwidth BW is inversely proportional to the resonance Q -factor: $\text{BW} = f_0/Q$, where f_0 is the resonance frequency. The Q -factor is composed of an intrinsic Q -factor Q_{in} and an input and output coupling induced external Q -factor Q_{ex} : $1/Q = 1/Q_{\text{in}} + 1/Q_{\text{ex}}$. Hence, small α (high Q_{in}) and small κ (high Q_{ex}) tend to have large Q and small resonance bandwidth. In reality, α is mainly determined by the fabrication process, but κ can be controlled by proper design of the gap between the straight waveguides and the microring resonator or by using directional couplers.

The basic building blocks are interconnected to form higher-order crossbars. In an $N \times N$ crossbar, each optical routing path between the input and output ports has N or $N - 1$ sequentially cascaded basic building blocks. As these basic building blocks are located in different columns, the microrings have different resonance wavelengths, and we denote their through and drop transmission functions as T_{tl} and T_{dl} (l is the column number). The crossbar output electric field transmission function T_{ij} ($i, j = 1, 2, \dots, N$ for input and output port numbers) is therefore the product

Fig. 2 Contour maps of (a) extinction ratio and (b) resonance 3-dB bandwidth as a function of ring waveguide propagation loss α and coupling coefficient κ

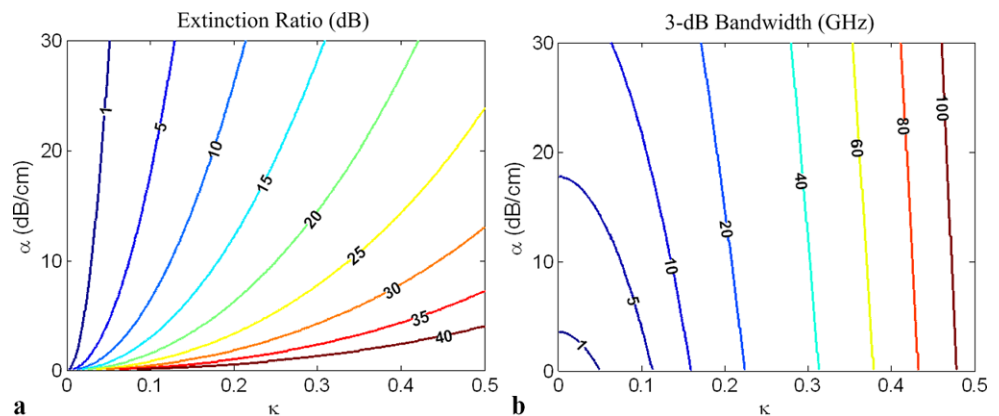


Fig. 3 Calculated transmission spectra for a 3×3 crossbar with various resonance bandwidths: (a) 200 GHz, (b) 50 GHz, and (c) 20 GHz. The channel spacing is fixed at 50 GHz


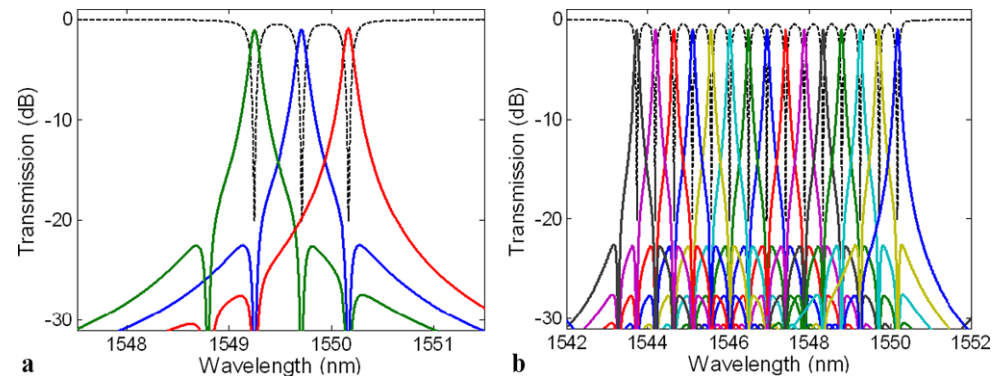


Fig. 4 Calculated transmission spectra for (a) a 4×4 crossbar and (b) a 16×16 crossbar, assuming a resonance bandwidth of 20 GHz and a channel spacing of 50 GHz



of several transmission functions T_{il} and T_{dl} . For instance, in the 3×3 crossbar, $T_{11} = T_{i1}T_{d2}T_{i3}$, $T_{12} = T_{d1}T_{i3}$, and $T_{13} = T_{i1}T_{i2}$.


In order to see the effect of cascading multiple microring resonators on the transmission performance, Fig. 3 illustrates the output spectra with various resonance Q -factors. The size of each column of resonators (ring radius $\sim 10 \mu\text{m}$) is gradually reduced such that the corresponding resonance is blue-shifted by 0.4 nm (~ 50 GHz). The resonance ER is maintained at 20 dB. Figures 3(a)–(c) show the output port transmission spectra (input from input port 1) in a 3×3 crossbar with $Q = 10^3$ (bandwidth ~ 200 GHz), 4×10^3 (~ 50 GHz), and 10^4 (~ 20 GHz), respectively. Figure 3(a) shows that, for low- Q resonances, the resonance line-shapes broaden and interfere more strongly with each other. As a result, the O_3 spectrum has deeper resonance dips, yet the desired resonance peak in the O_1 spectrum is heavily suppressed by the two adjacent resonances. As Fig. 3(b) illustrates, the high Q values allow the corresponding bandwidth to approach the channel spacing, and as a result, the O_1 and O_2 transmissions are reasonably clear with channel crosstalk < -10 dB. Figure 3(c) shows the case where further increases in Q offer sharper resonance line-shapes with lower optical crosstalk. It should be noted that the resonance bandwidth should be large enough to be able to support the optical signal modulation bandwidth in order not to cause signal distortion during transmission.

Figures 4(a) and (b) show simulated output transmission spectra for 4×4 and 16×16 crossbars. The channel bandwidth is 20 GHz, the channel spacing 50 GHz, and the resulting crosstalk -20 dB. Note that passband shapes with broader bandwidths and sharper roll-offs can be achieved by using series-coupled microring resonators instead of the single microring resonator in the basic building block.

4 Fabrication and measurements

We fabricated our devices using 6 inch silicon-on-insulator (SOI) wafers with top silicon layer thicknesses of 200 and 260 nm. The buried oxide (BOX) layer thickness is 2 μm . The wafers were coated with Shipley UV-210 photoresist at 7000 rpm with a thickness of 0.42 μm , and then soft-baked at 130°C for 1 minute. The patterns were exposed by 248-nm Deep-UV (DUV) light using an ASML 5500/90 stepper. The exposed wafers were post-exposure baked for 1 minute and then developed in Shipley LDD 26W. UV-baking of the patterned photoresist enhanced its etch selectivity in favor of silicon etching during the subsequent HBr-Cl₂ gas-based reactive-ion-etching (RIE) in a Transformer Coupled Plasma (TCP) Lam Etcher. The devices were finally deposited with low temperature oxide (LTO) in a low pressure chemical vapor deposition (LPCVD) furnace. The oxide upper-cladding

Fig. 5 Scanning electron microscope (SEM) images of (a) a single microring resonator, (b) waveguide cross-section before oxide deposition, and (c) waveguide cross-section after oxide deposition



layer can enhance the coupling between the waveguides and the microring resonator, while passivating the etched sidewalls of the silicon waveguides.

Figure 5(a) shows a top-view scanning electron microscope (SEM) image of a fabricated single microring resonator. The microring resonator radius is $10\ \mu\text{m}$. The waveguide width is $\sim 0.45\ \mu\text{m}$ and the height is $\sim 0.2\ \mu\text{m}$ with a $\sim 40\ \text{nm}$ silicon layer left as the slab region. The gap between the waveguide and microring resonator is $0.25\ \mu\text{m}$. Figure 5(b) shows the waveguide cross-section before LTO deposition. Slight roughness on the waveguide sidewall is visible. Figure 5(c) shows the waveguide cross-section after LTO deposition ($\sim 0.7\ \mu\text{m}$ thick). The wafer was subsequently diced into mm-sized dies and the die facets were polished for high-quality optical coupling.

The device transmission spectra measurement employed an external-cavity wavelength-tunable diode laser ($1524\text{--}1576\ \text{nm}$, $\sim 300\ \text{kHz}$ linewidth) as the light source. The input laser beam was first pre-amplified by an erbium doped fiber amplifier (EDFA) to $\sim 16\ \text{dBm}$ optical power to compensate for the device insertion loss ($\sim 20\ \text{dB}$). We used a polarization controller to set the polarization to TE polarization (electric field parallel to the chip). The input light was butt-coupled to the device through a tapered single-mode silica fiber. The tapered fiber has a spot size of $\sim 2.5\ \mu\text{m}$. The input waveguides of the devices are also laterally tapered to $\sim 2.5\ \mu\text{m}$ to reduce coupling loss. At the output, we used an objective lens with numerical aperture $\text{NA} = 0.65$ to collect the output light. The collected light was finally detected by an InGaAs photodiode detector.

Figures 6(a) and (b) show the crossbar basic building block microscope image and its measured through and drop transmission spectra. The resonance free spectral range (FSR) is $\sim 14\ \text{nm}$ and the resonance bandwidth is $\sim 1.6\ \text{nm}$ ($200\ \text{GHz}$). Figures 6(c) and (d) show the fabricated 3×3 crossbar and the corresponding measured spectra for its three output ports. The resonance wavelengths of the three resonators are separated by $\sim 75\ \text{GHz}$. As predicted

by the modeling, when the resonance bandwidth is larger than the channel spacing, strong interference occurs among resonators, and consequently the drop transmission spectra (O_1 and O_2) exhibit less pronounced peaks. These weakened drop peaks deteriorate optical signal transmission. Figure 6(e) shows the fabricated 4×4 crossbar using the SOI wafers with a device layer thickness of $0.26\ \mu\text{m}$. Figure 6(f) shows the corresponding measured spectra for its four output ports. As the light transmission goes through more stages, cross interference between different resonators is more significant, which increases the channel crosstalk.

Fabrication-induced phase error in the microring resonators can cause resonance shift from the desired values. In order to compensate for this error, post-fabrication trimming techniques, such as electron beam trimming [23], can be used to accurately control the resonance wavelengths. Alternatively, doping the desired waveguide regions with p-type or n-type dopants (e.g., boron and phosphorus) is also able to change the refractive index [24].

We characterized the 3×3 crossbar dynamic performance by measuring the transmission bit error rate (BER) at 10 and 40 Gbit/s data rates. The pseudo-random binary sequence (PRBS) is $2^{31} - 1$ bits long for the 10 Gbit/s measurement and $2^7 - 1$ bits long for the 40 Gbit/s measurement. Figure 7 shows the setup for the BER measurement. 10 and 40 Gbit/s NRZ-OOK optical signals were generated by modulating an input laser beam using a Lithium Niobate (LiNbO_3) Mach-Zehnder modulator driven by a pulse pattern generator (PPG) at 10 and 40 GHz. The modulated signal was then amplified by an EDFA (saturated at 16 dBm optical power) before coupling into the chip. The transmitted signal was post-amplified by another EDFA (saturated at 16 dBm optical power) and then passed through a tunable band pass filter (BPF). The optical signal was finally converted to an electrical RF signal after a high speed photodetector, and analyzed by a bit error rate tester (BERT), which is synchronized to the PPG.

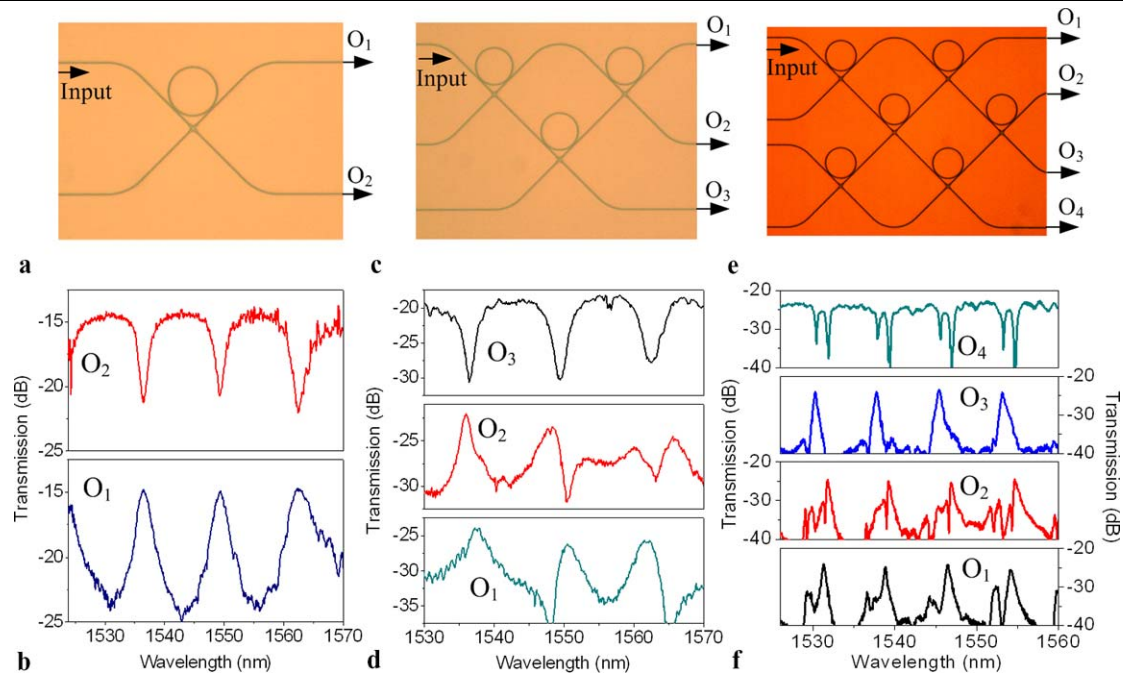
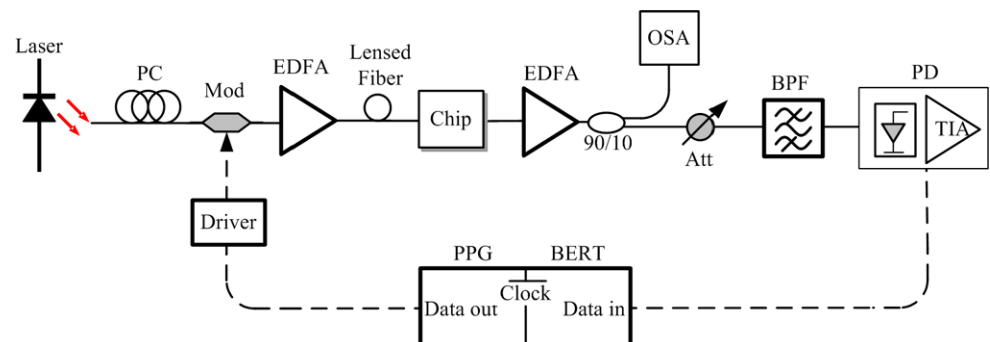


Fig. 6 (a) Microscope top view of the fabricated crossbar basic building block, (b) measured transmission spectra for the device in (a), (c) microscope top view of the fabricated 3×3 crossbar, (d) mea-

sured transmission spectra of the device in (c). (e) Microscope top view of the fabricated 4×4 crossbar, (f) measured transmission spectra of the device in (e)

Fig. 7 Experimental setup for measuring the bit-error rate. PC: polarization controller, Mod: modulator, EDFA: erbium-doped fiber amplifier, OSA: optical spectrum analyzer, Att: attenuator, BPF: band-pass filter, PD: photodiode, PPG: pulse pattern generator, BERT: bit error rate tester



We tuned the laser wavelength to the resonance wavelengths (near 1550 nm) for O_1 and O_2 BER measurements and a non-resonance wavelength (near 1543 nm) for the O_3 BER measurement. Figures 8(a) and (b) show the BER test results for all three output ports and back-to-back (BtB) transmission at 10 Gbit/s and 40 Gbit/s. The power penalty for O_3 is smaller than that for O_1 and O_2 as expected, since the signal coming out of O_3 travels through the crossbar without experiencing any resonance. The power penalties for the drop channels are slightly higher, resulting from the limited resonance bandwidth and interference induced resonance line-shape distortion (Fig. 6(d)) [25, 26].

The highest order of the crossbar is essentially determined by the ratio of the resonance FSR and the channel spacing. Smaller microrings have large FSRs, yet they also suffer more bending losses. The smallest microring resonator that has been experimentally demonstrated is 3 μm in

diameter, with a FSR of ~ 63 nm [11]. As discussed previously, in order for the WDM signals transmission through the crossbar without significant deterioration, the WDM channel spacing should be larger than the resonance bandwidth, and the resonance bandwidth should be larger than the signal bandwidth. Thus, for 40 Gbit/s WDM signals, the channel spacing can be 0.64 nm (double the resonance bandwidth) to tolerant ± 0.1 nm fabrication error-induced resonance shift with a crosstalk of < -10 dB. In this case, one resonance FSR can incorporate ~ 100 WDM channels, and 100×100 crossbar can be built to cross-connect 100 sources and destinations. Each channel suffers ~ 10 dB loss due to the ~ 100 stages of waveguide crossings. We believe the 100×100 crossbar can be realized using current fabrication technologies, with careful trimming of each microring resonator.

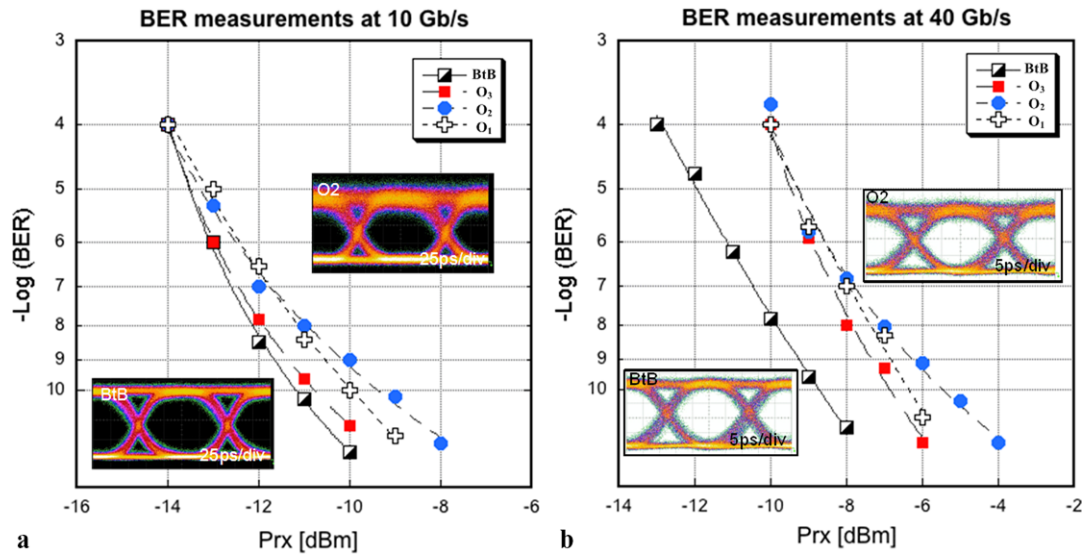


Fig. 8 Bit-error rate curves for the setup back-to-back transmission (BtB) and three output-ports (O_1 , O_2 and O_3) at (a) 10 Gbit/s and (b) 40 Gbit/s data rates. Insets show the eye-diagrams for BtB and O_2 transmissions




Fig. 9 Schematic diagrams of (a) single-to-multiple transmission and (b) multiple-to-single transmission

5 Applications in future multicore processors

In addition to functioning as a one-to-one cross-connect network, our crossbar also has the unique ability to support many-to-one (multiplexing) and one-to-many (multicasting) transmissions that arise in many applications, without additional cost or complexity. This opens up opportunities for the proposed crossbar design in future multicore processors.

The single-to-multiple transmission, or multicast [27], is the transmission from a single source to multiple destinations simultaneously. Consider the simple 4×4 crossbar shown in Fig. 9(a). A network node connected to input I_1 sends four packets using four different wavelengths simultaneously ($\lambda_1, \lambda_2, \lambda_3$, and λ_4) to destination nodes connected to O_1, O_2, O_3 , and O_4 , respectively. The packet on λ_1 is diverted at the first microring resonator and then forwarded to O_3 . The other three packets do the same to their destinations, which results in the multicasting of a single packet to four destinations simultaneously. Let us consider implementing cache coherency in a multicore processor. In both

snooping and directory-based coherence protocols, write-invalidate messages have to be sent to multiple destinations when a shared cache block is updated. In [27], the authors demonstrate that even a small fraction of multicast traffic can significantly degrade the overall performance and hence argue for hardware support of multicasting in future on-chip networks. We believe the proposed crossbar design can be useful in such applications without additional cost or complexity and in fact is more efficient than the virtual-circuit minimum spanning tree approach proposed by the authors in [27] to improve the performance of cache-coherence protocols.

The ability to send data from multiple sources to a single destination simultaneously is also very useful in many applications—not only in realizing scalable implementations of cache coherence protocols, but also in architectures that are derivatives of the dataflow model of computation such as RAW [28, 29], WaveScalar [30] and TRIPS [31]. In a typical interconnection network, this is achieved by multiple sources arbitrating for the right to communicate with the

single destination, which involves serialization at the switch fabric level. With the proposed crossbar, the destinations can accept the packets in multiple wavelengths simultaneously and separate them into different packets using tunable filters as shown in Fig. 9(b).

6 Summary

Large scale low latency crossbars are very desirable in on-chip interconnect networks. However, crossbars with a large number of ports and wide data paths are not practical in an electronic implementation because of wiring complexity and the large and power hungry drivers and repeaters needed to support low latency operation. We propose and demonstrate a microring resonator-based arbitration-free optical crossbar that can be used for on-chip WDM based interconnection networks. The proposed crossbar design supports multicasting and multiplexing efficiently which can be useful in many applications in future manycore processors. The crossbar design is modeled using the transfer matrix method to study the effect of several key parameters on its transmission performance. Experimental results indicate error-free operation of a 3×3 crossbar up to 40 Gbit/s.

Acknowledgements This project was supported in part by the Center for Information Technology in the Interest of Society (CITRIS) seed funds # 3-34PICOB.

Open Access This article is distributed under the terms of the Creative Commons Attribution Noncommercial License which permits any noncommercial use, distribution, and reproduction in any medium, provided the original author(s) and source are credited.

References

1. D.A.B. Miller, Proc. IEEE **88**, 728 (2000)
2. P.R.A. Binetti, X.J.M. Leijtens, M. Nikoufard, R. Orobtschouk, T. Benyattou, T. de Vries, Y.S. Oei, M.K. Smit, in *Proceedings of IEEE/LEOS Symposium (Benelux Chapter)*, Mons, Belgium (2005), p. 233
3. H.H. Chang, A.W. Fang, M.N. Sysak, H. Park, R. Jones, O. Cohen, O. Raday, M.J. Paniccia, J.E. Bowers, Opt. Express **15**, 11466 (2007)
4. L. Liao, D. Samara-Rubio, M. Morse, A. Liu, D. Hodge, D. Rubin, U. Keil, T. Franck, Opt. Express **13**, 3129 (2005)
5. A. Liu, R. Jones, L. Liao, D. Samara-Rubio, D. Rubin, O. Cohen, R. Nicolaescu, M. Paniccia, Nature **427**, 615 (2004)
6. A. Liu, L. Liao, D. Rubin, H. Nguyen, B. Ciftcioglu, Y. Chetrit, N. Izhaky, M. Paniccia, Opt. Express **15**, 660 (2007)
7. J. Liu, J. Michel, W. Giziewicz, D. Pan, K. Wada, D.D. Cannon, S. Jongthammanurak, D.T. Danielson, L.C. Kimerling, J. Chen, Appl. Phys. Lett. **87**, 103501 (2005)
8. M.S. Nawrocka, T. Liu, X. Wang, R.R. Panepucci, Appl. Phys. Lett. **89**, 071110 (2006)
9. J. Van Campenhout, P. Rojo Romeo, P. Regreny, C. Seassal, D. Van Thourhout, S. Verstuyft, L. Di Cioccio, J.M. Fedeli, C. Lagae, R. Baets, Opt. Express **15**, 6744 (2007)
10. S. Xiao, M.H. Khan, H. Shen, M. Qi, J. Lightwave Technol. **26**, 228 (2008)
11. Q. Xu, D. Fattal, R.G. Beausoleil, Opt. Express **16**, 4309 (2008)
12. Q. Xu, S. Manipatruni, B. Schmidt, J. Shakya, M. Lipson, Opt. Express **15**, 430 (2007)
13. Q. Xu, B. Schmidt, J. Shakya, M. Lipson, Opt. Express **14**, 9431 (2006)
14. T. Yin, R. Cohen, M.M. Morse, G. Sarid, Y. Chetrit, D. Rubin, M.J. Paniccia, Opt. Express **15**, 13965 (2007)
15. W.M. Green, M.J. Rooks, L. Sekaric, Y.A. Vlasov, Opt. Express **15**, 17106 (2007)
16. S. Manipatruni, Q. Xu, M. Lipson, Opt. Express **15**, 13035 (2007)
17. R.G. Beausoleil, P.J. Kuekes, G.S. Snider, S.Y. Wang, R.S. Williams, Proc. IEEE **96**, 230 (2008)
18. A. Shacham, K. Bergman, L.P. Carloni, in *Proceedings of 44th ACM/IEEE Design Automation Conference (DAC '07)*, ed. by K. Bergman (2007), p. 132
19. B.E. Little, S.T. Chu, W. Pan, Y. Kokubun, IEEE Photonics Technol. Lett. **12**, 323 (2000)
20. H. Chen, A.W. Poon, IEEE Photonics Technol. Lett. **18**, 2260 (2006)
21. L. Zhang, M. Yang, Y. Jiang, E. Regentova, E. Lu, in *Proceedings of the 18th IASTED International Conference on Parallel and Distributed Computing and System*, Dallas, TX, 2006
22. A. Yariv, Electron. Lett. **36**, 321 (2000)
23. J. Schrauwen, D. Van Thourhout, R. Baets, Opt. Express **16**, 3738 (2008)
24. R. Soref, B. Bennett, IEEE J. Quantum Electron. **23**, 123 (1987)
25. B.G. Lee, B.A. Small, K. Bergman, Q. Xu, M. Lipson, Opt. Lett. **31**, 2701 (2006)
26. B.G. Lee, B.A. Small, Q. Xu, M. Lipson, K. Bergman, IEEE Photonics Technol. Lett. **19**, 456 (2007)
27. N.E. Jerger, L.S. Peh, M. Lipasti, in *Proceedings of 35th International Symposium on Computer Architecture (ISCA'08)* (2008), p. 229
28. S. Swanson, K. Michelson, A. Schwerin, M. Oskin, in *Proceedings of the 36th Annual IEEE/ACM International Symposium on Microarchitecture (MICRO-36)* (2003), p. 291
29. A. Putnam, S. Swanson, M. Mercaldi, K. Michelson, A. Petersen, A. Schwerin, M. Oskin, S. Eggers, Technical Report UWCSE-2005-10-2, UW-Computer Science and Engineering, 2005
30. M.B. Taylor, W. Lee, S.P. Amarasinghe, A. Agarwal, IEEE Trans. Parallel Distrib. Syst. **16**, 145 (2005)
31. K. Sankaralingam, R. Nagarajan, H. Liu, C. Kim, J. Huh, D. Burger, S.W. Keckler, C.R. Moore, ACM SIGARCH Comput. Archit. News **31**, 422 (2003)