

Scalability of Optical Interconnects Based on Microring Resonators

Andrea Bianco, Davide Cuda, Roberto Gaudino, Guido Gavilanes, Fabio Neri, and Michele Petracca

Abstract—This letter investigates the use of optical microring resonators as switching elements (SEs) in large optical interconnection fabrics. We introduce a simple physical-layer model to assess scalability in crossbar- and Benes-based architectures. We also propose a new dilated SE that improves scalability to build fabrics of several terabits per second of aggregate capacity.

Index Terms—Microring resonators (MRs), optical interconnects.

I. INTRODUCTION

INTEGRATED electronic interconnections are largely used in commercial packet-switching on-chip architectures, where high-speed switching fabrics are required. Predictions in the International Technology Roadmap for Semiconductors show that latency and power requirements of silicon electrical interconnections (for wiring lengths above the millimeter) limit the performance of electrical on-chip interconnections. Recent technological breakthroughs in silicon photonic integration offer several solutions to overcome these limitations in the medium to long term.

Silicon *microring resonators* (MRs) have been recently studied in different application domains. We investigate MRs' capability to build switching elements (SEs) to design large integrated optical switching fabrics, used to interconnect external optical transceivers in a board-to-board scenario. We focus exclusively on the design of photonic fabrics, without considering electrooptical (E-O) conversion issues. We limit our analysis to single wavelength operation: all transmitters (TXs) use, and all rings resonate at the same wavelength. The analysis can be extended to a wavelength-division-multiplexing (WDM) scenario, by assuming wavelength striping, with WDM channels fitting exactly the periodical MRs' transfer function [1].

We propose a physical-layer model of different microring-based SEs, building upon the experimental results reported in [1]. Then, we assess the scalability of MR-based crossbar and Benes architectures to build optical fabrics. Finally, we propose

Manuscript received January 14, 2010; revised March 12, 2010; accepted April 25, 2010. Date of publication May 17, 2010; date of current version June 30, 2010. This work was supported by the BONE project, a NoE funded by the European Commission within the 7th Framework Programme.

A. Bianco, D. Cuda, R. Gaudino, G. Gavilanes, and F. Neri are with Politecnico di Torino, Torino 10129, Italy (e-mail: andrea.bianco@polito.it; davide.cuda@polito.it; roberto.gaudino@polito.it; guido.gavilanes@polito.it; fabio.neri@polito.it).

M. Petracca is with Columbia University, New York, NY 10027 USA (e-mail: petracca@cs.columbia.edu).

Digital Object Identifier 10.1109/LPT.2010.2049834




Fig. 1. 1B-SE in the (a) tuned and (b) untuned states.

a dilated SE architecture, that can boost the fabric aggregate capacity to several terabits per second.

II. MICRORING MODEL BASICS

MRs are based on a circular waveguide coupled to one or two straight waveguides. Fig. 1 shows a simple configuration to build a basic 1×2 SE, labeled 1B-SE.

In the 1B-SE, an optical signal entering the *input* port can be deflected either to the *drop* port (when the ring is tuned to the signal wavelength), or to the *through* port (if the ring is untuned). The ring resonating wavelength can be dynamically changed (i.e., tuned) by properly adjusting some component parameter, such as the ring effective refractive index. If the latter is changed by carrier injection using p-i-n junctions with an external electrical signal, the switching time can be very fast (in the ns range as shown in [2]). This property, combined with a very small footprint, makes these structures very attractive for large-scale integrated photonic switching fabrics.

We develop a simple transmission model for the 1B-SE in terms of optical attenuation and crosstalk, as shown in Fig. 1. As a convention, all formulas are in linear units, while numerical values and results are in decibels.

- 1) The input signal suffers an insertion loss IL_d^{on} when the ring is in the *drop* configuration [Fig. 1(a)], and IL_t^{on} when in the *through* configuration [Fig. 1(b)].
- 2) In both configurations, a residual optical power appears on the nominally unused outputs, characterized by two insertion losses IL_d^{off} and IL_t^{off} .

Ideally, $IL_d^{\text{on}} = IL_t^{\text{on}} = 0$ dB, and $IL_d^{\text{off}} = IL_t^{\text{off}} = 0$ dB are extremely high. In practice, as measured in [1], $IL_d^{\text{on}} \simeq 1.4$ dB, $IL_d^{\text{off}} \simeq 18.1$ dB, and $IL_t^{\text{off}} \simeq 23.1$ dB; we assume $IL_t^{\text{on}} \simeq 0.1$ dB. Thus, the 1B-SE structure shows a largely asymmetric behavior, i.e., lower losses in the untuned (no turn) case. Finally, we also define the extinction ratio (ER) between the ON and OFF states of each port as $ER_d = IL_d^{\text{off}}/IL_d^{\text{on}}$ and $ER_t = IL_t^{\text{off}}/IL_t^{\text{on}}$, respectively.

Classical interconnects require 2×2 SEs as building blocks: we introduce in Fig. 2(a) a 2×2 basic SE (2B-SE), an SE experimentally demonstrated in [3]. Two 1B-SEs jointly controlled provide two possible states, indicated as bar(In1 \rightarrow Out1, In2 \rightarrow




Fig. 2. 2×2 SEs in the 2B-SE and 2D-SE configurations. (a) 2B-SE; (b) 2D-SE.

Out2) and cross(In1 \rightarrow Out2, In2 \rightarrow Out1). Also the 2B-SE exhibits an asymmetric behavior. Denoting the insertion losses in the bar and cross states as $\mathbb{I}L^{\text{bar}}$ and $\mathbb{I}L^{\text{cross}}$, $\mathbb{I}L^{\text{bar}} = \mathbb{I}L_d^{\text{on}}$ and $\mathbb{I}L^{\text{cross}} = (\mathbb{I}L_t^{\text{on}})^2$. We also characterize the crosstalk by introducing the parameter $X_{2\text{B-SE}} = P_{\text{xtalk}}/P_{\text{useful}}$ for each output port. Assuming that the optical powers at the two input ports are nominally equal

$$X_{2\text{B-SE}}^{\text{bar}} = \frac{\mathbb{I}L_d^{\text{on}}}{(\mathbb{I}L_t^{\text{off}})^2} \quad \text{and} \quad X_{2\text{B-SE}}^{\text{cross}} = \frac{(\mathbb{I}L_t^{\text{on}})^2}{\mathbb{I}L_d^{\text{off}}}. \quad (1)$$

The 2×2 SE shows an asymmetric behavior because $X_{2\text{B-SE}}^{\text{bar}} < X_{2\text{B-SE}}^{\text{cross}}$ (-44.6 dB versus -17.8 dB).

To reduce asymmetries of microring-based SEs, we introduce the 2×2 dilated SE (2D-SE) depicted in Fig. 2(b). Different from 1B-SEs and 2B-SEs, the 2D-SE presents the same power penalties regardless of the SE state. Indeed, the incoming signals are always deflected exactly once. In more detail, in the cross (bar) state, the signal at In1 enters (goes straight) the first ring and goes straight (enters) in the second ring. The same holds for the signal entering at In2. The total loss for the useful signal is $\mathbb{I}L_{2\text{D-SE}} = \mathbb{I}L_t^{\text{on}} \cdot \mathbb{I}L_d^{\text{on}}$ on both the bar and cross states. Besides obtaining this symmetry, the 2D-SE forces leakage signals to pass two rings with loss $\mathbb{I}L_t^{\text{off}} \cdot \mathbb{I}L_d^{\text{off}}$, resulting in crosstalk levels of 39.7 dB

$$X_{2\text{D-SE}}^{\text{cross}} = X_{2\text{D-SE}}^{\text{bar}} = \frac{\mathbb{I}L_d^{\text{on}} \cdot \mathbb{I}L_t^{\text{on}}}{\mathbb{I}L_d^{\text{off}} \cdot \mathbb{I}L_t^{\text{off}}} = \frac{1}{\text{ER}_d \cdot \text{ER}_t}. \quad (2)$$

III. INTERCONNECTION ARCHITECTURES

We study three silicon photonic interconnections to build an on-chip fabric for board-to-board interconnections. Benes networks (in two flavors, respectively using 2B-SEs or 2D-SEs) are chosen because they exhibit, among nonblocking architectures, the minimum complexity, measured as the number of needed MRs. Instead, the crossbar minimizes the transmission impairments, improving scalability, which is assessed calculating the signal degradation along the worst signal path.

A. Scalability of Microring-Based Crossbars

An $N \times N$ crossbar can be implemented as a matrix of N^2 1×2 SEs [Fig. 3(a)]. This architecture exhibits good scalability, because each input signal can reach any output port crossing a single tuned MR, at the expenses of a high complexity, equal to $C_{\text{xbar}}(N) = N^2$. More in detail, the worst-case path is the one connecting input In1 to output OutN, which gives a useful output power $P_{\text{out}} = P_{\text{in}}/\mathbb{I}L_{\text{xbar}}$, where $\mathbb{I}L_{\text{xbar}} = \mathbb{I}L_d^{\text{on}} \cdot (\mathbb{I}L_t^{\text{on}})^{2(N-1)}$. For each of the $N-1$ remaining




Fig. 3. Interconnection architectures and critical signal paths. (a) Crossbar; (b) Benes network.

inputs, a crosstalk signal reaches the bottom row, as shown in Fig. 3(a). Before reaching the last ring of the column, each crosstalk signal has power $P_{\text{CT}} = P_{\text{in}}/(\mathbb{I}L_t^{\text{off}} \cdot (\mathbb{I}L_t^{\text{on}})^{N-2})$, that accounts for one tuned and $N-2$ untuned 1B-SEs. Then, each crosstalk signal exits from the drop port of the last row of rings, and flows on to the right, through all the remaining rings on the last row; hence, the total crosstalk power becomes $P_{\text{xtalk}} = (P_{\text{CT}}/\mathbb{I}L_d^{\text{off}}) \sum_{k=0}^{N-2} (1/\mathbb{I}L_t^{\text{on}})^k$. Finally, normalizing the crosstalk to the useful output power P_{out} , and recalling ER_d and ER_t definitions, we have

$$X_{\text{xbar}} = \frac{(\mathbb{I}L_t^{\text{on}})^{N-1}}{\text{ER}_d \cdot \text{ER}_t} \sum_{k=0}^{N-2} \left(\frac{1}{\mathbb{I}L_t^{\text{on}}} \right)^k. \quad (3)$$

B. Scalability of Microring-Based Benes

Multistage $N \times N$ Benes networks show a much lower complexity than crossbars: $C_{\text{Benes}}(N) = (N/2) \cdot C_{\text{SE}} \cdot S(N)$, where $S(N) = 2 \log_2(N) - 1$ is the number of stages (e.g., in Fig. 3(b), $N = 8$ and $S(8) = 5$), and C_{SE} is the complexity of the single 2×2 SE used. Thus, $C_{\text{SE}} = 2$ or $C_{\text{SE}} = 4$ for the 2B-SE or the 2D-SE, respectively. The scalability of a Benes network depends on the number of network stages optical signals cross from inputs to outputs. When using 2B-SEs, where $\mathbb{I}L^{\text{bar}} > \mathbb{I}L^{\text{cross}}$ and $X^{\text{bar}} < X^{\text{cross}}$, we distinguish among the loss worst-case path (LWCP) and the crosstalk worst-case path (XWCP), as indicated in Fig. 3(b) with continuous and dashed lines, respectively. Signals suffer an LWCP when they cross SEs all configured in the bar state. Each bar-state element at the j th ($1 \leq j \leq S(N)$) stage has an output power equal to $P_j = P_{j-1}/\mathbb{I}L^{\text{bar}}$; thus, after $S(N)$ stages, $P_{S(N)} = P_{\text{in}}/(\mathbb{I}L^{\text{bar}})^{S(N)}$. Conversely, the XWCP occurs when input signals go through cross-state SEs only. At stage j , the output crosstalk power is $X_j = X_{j-1} + X^{\text{cross}}$ (considering the worst case in which the powers at the two input ports are equal). After $S(N)$ stages

$$X_{\text{Benes}}(N) = S(N) \cdot X^{\text{cross}}. \quad (4)$$

Note that (4) holds when either 2B-SEs or 2D-SEs are employed as building SEs, with $X^{\text{cross}} = X_{2\text{B-SE}}^{\text{cross}}$ or $X^{\text{cross}} = X_{2\text{D-SE}}^{\text{bar,cross}}$ respectively.

IV. SCALABILITY RESULTS

To assess the scalability of the microring-based interconnection networks, we employed the coherent crosstalk penalty model presented in [4], that estimates the additional power




Fig. 4. Power penalties for all fabrics.

required at the receiver in presence of coherent crosstalk to maintain a target bit-error rate (T_{BER}) as

$$IX = \frac{1}{(1 - X \cdot Q^2)} \quad (5)$$

where X is the corresponding total normalized crosstalk power and Q is the quality factor set by T_{BER} : for instance, for $T_{BER} = 10^{-12}$, $Q = 7$, assuming a nonreturn-to-zero (NRZ) modulation, optimum sampling times, and negligible electrical noise. To assess scalability at different bit rates (R_b), we used the receiver sensitivity model presented in [5], that assumes a sensitivity ($P_S(R_b)$) slope versus R_b of 13.5 dB/decade and a reference value of -26 dBm at 10 Gb/s. For the transmitter, a typical average transmitted power $P_{TX} = 3$ dBm is assumed, accounting for modulation, ON-OFF keying and coupling to the fabric. Finally, the received power P_{RX} is conditioned to the system power budget

$$P_{RX} = \frac{P_{TX}}{IL \cdot IX \cdot \mu} \geq P_S(R_b) \quad (6)$$

where IL and IX are the total losses and crosstalk penalties, and μ is a system margin, conservatively set to 3 dB to account for component aging and other non modeled effects. For the 1B-SE, we set $ER_d = 16.7$ dB, $ER_t = 23$ dB, $IL_d^{on} = 1.4$ dB, and $IL_t^{on} = 0.1$ dB, as measured in [1]. The waveguide crossing loss is already taken into account by the value IL_t^{on} . For Benes crossing, we assumed 0.04 dB loss per crossing [6].

Scalability is limited either by losses or crosstalk penalties. Fig. 4 reports crosstalk and loss penalties versus the total number of ports N . To check the analytical formulation, we performed a time-domain Monte Carlo simulation (TDMCS) to estimate the crosstalk penalty in the crossbar IX case. Simulation results show a good agreement with the theory. The crosstalk penalty in the crossbar is the result of the aggregation of many different crosstalk sources (along the XWCP); this source aggregation was simulated and the impact of crosstalk in transmission was quantified. Crosstalk limits scalability for both the 2B-SE Benes and crossbar architectures. Even a single stage in 2B-SE Benes networks presents such a high crosstalk that generates unacceptable penalties. 2D-SEs drastically reduce crosstalk penalties, thanks to an improved SE extinction ratio, at the expenses of higher losses on the received signal. Fig. 5 shows the maximum capacities achieved per wavelength and the required number of MRs, as a function of the single

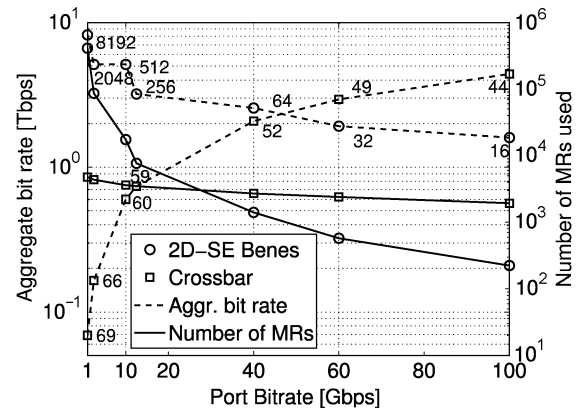


Fig. 5. Aggregate bit rate and complexity comparison.

transmitter bit rate. The total number of fabric ports N is also reported. The graph is obtained by inserting results shown in Fig. 4 with the receiver sensitivity constraint given in (6). Only feasible configurations are plotted.

At low bit rates (i.e., 1 Gb/s), 2D-SE Benes reach the highest aggregated bandwidth at about 8.1 Tb/s, because the receivers' sensitivity tolerates such a high insertion loss. Instead, crosstalk penalties prevail in the crossbar allowing sharply fewer ports (≈ 50), and much lower aggregate bandwidth. At high bit rates, high aggregate bandwidths can also be achieved: e.g., 4.4 Tb/s for the crossbar and 1.6 Tb/s for the 2D-SE Benes with 100-Gb/s receivers. However, the crossbar requires 10 times more rings than 2D-SE Benes. Amplification stages, not considered here, can boost 2D-SE Benes performance, because penalties are mainly induced by insertion losses, both at low and high bit rates, while the crosstalk is not the limiting factor yet (Fig. 4). Note that the robustness to crosstalk of the Benes topology is given by the ring redundancy of the 2D-SE, which, however, increases the attenuation due to the higher number of used rings (4 instead of 2 in a 2B-SE).

Although MR fabrication is still in an embryonic phase, we believe that we showed that MRs are a promising candidate to support terabits per second switching matrices in on-chip interconnections.

REFERENCES

- [1] B. G. Lee, A. Biberman, P. Dong, M. Lipson, and K. Bergman, "All-optical comb switch for multiwavelength message routing in silicon photonic networks," *IEEE Photon. Technol. Lett.*, vol. 20, no. 10, pp. 767-769, May 15, 2008.
- [2] C. Li and A. W. Poon, "Silicon electro-optic switching based on coupled-microring resonators," in *Proc. Conf. Lasers and Electro-Optics (CLEO)*, Baltimore, MD, May 2007.
- [3] B. G. Lee, A. Biberman, N. Sherwood-Droz, C. B. Poitras, M. Lipson, and K. Bergman, "High-speed 2×2 switch for multi-wavelength message routing in on-chip silicon photonic networks," in *Proc. ECOC*, Brussels, Belgium, Sep. 2008.
- [4] H. Takahashi, K. Oda, and H. Toba, "Impact of crosstalk in an arrayed-waveguide multiplexer on $N \times N$ optical interconnection," *J. Lightw. Technol.*, vol. 14, no. 6, pp. 1097-1105, Jun. 1996.
- [5] E. Sackinger, *Broadband Circuits for Optical Fiber Communication*. New York: Wiley, 2005.
- [6] H. Bukkems, C. Herben, M. Smit, F. Groen, and I. Moerman, "Minimization of the loss of intersecting waveguides in InP-based photonic integrated circuits," *IEEE Photon. Technol. Lett.*, vol. 11, no. 11, pp. 1420-1422, Nov. 1999.
- [7] R. Gaudino, G. A. G. Castillo, F. Neri, and J. M. Finochietto, "Simple optical fabrics for scalable terabit packet switches," in *Proc. IEEE ICC*, Beijing, China, May 2008.