

Scalability of Optical Interconnects Based on Microring Resonators

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

Abstract—This paper investigates the use of optical microring resonators as switching elements 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 switching element that improves scalability to build fabrics of several Tbps of aggregate capacity.

Index Terms—Optical interconnects, Microring resonators.

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 mm) 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 (MR) 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 Electro-Optical (E-O) conversion issues.**

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 MRs-based crossbar and Benes architectures to build optical fabrics. Finally, we propose a dilated SE architecture, that can boost the fabric aggregate capacity to several Tbps.

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, labelled 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

This work was supported by the BONE project, a NoE funded by the European Commission within the 7th Framework Programme. All authors are with Dip. di Elettronica, Politecnico di Torino, Italy, Email: {firstname.lastname}@polito.it except M. Petracca, which is with Columbia University, New York, USA. Email: petracca@cs.columbia.edu




Fig. 1: 1B-SE in the tuned and untuned states

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 dB.

- 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))
- 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, $IL_d^{on} = IL_t^{on} = 0$ dB, and $IL_d^{off} = IL_t^{off}$ are extremely high. In practice, as measured in [1], $IL_d^{on} \simeq 1.4$ dB, $IL_d^{off} \simeq 18.1$ dB, and $IL_t^{off} \simeq 23.1$ dB; we assume $IL_t^{on} \simeq 0.1$ 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 $ER_d = IL_d^{off} / IL_d^{on}$ and $ER_t = IL_t^{off} / IL_t^{on}$ respectively. Classical interconnects require 2×2 switching elements as building blocks: we introduce in Fig. 2(a) a 2×2 Basic SE (2B-SE), a switching element experimentally demonstrated in [3]. Two 1B-SEs jointly controlled provide two possible states, indi-




Fig. 2: 2×2 SEs in the 2B-SE and 2D-SE configurations

cated as *bar* ($in1 \rightarrow out1, in2 \rightarrow out2$) 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 IL^{bar} and IL^{cross} , $IL^{bar} = IL_d^{on}$ and $IL^{cross} = (IL_t^{on})^2$. We also characterize the crosstalk by introducing the parameter $X_{2B-SE} = P_{xtalk}/P_{useful}$ for each output port. Assuming that the optical powers at the two input ports are nominally equal:

$$X_{2B-SE}^{bar} = \frac{IL_d^{on}}{(IL_t^{off})^2} \quad \text{and} \quad X_{2B-SE}^{cross} = \frac{(IL_t^{on})^2}{IL_d^{off}} \quad (1)$$

The element shows an asymmetric behavior because $X_{2B-SE}^{bar} < X_{2B-SE}^{cross}$ (-44.6dB vs. -17.8dB): the crosstalk is smaller in the bar case.

To reduce asymmetries of microring-based SEs, we introduce the 2x2 Dilated SE (2D-SE) depicted in Fig. 2(b). Differently 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. More in details, in the cross (bar) state, the signal at *in1* enters (goes straight) the first ring and goes straight (enters) the second ring. The same holds for the signal entering at *in2*. The total loss for the useful signal is $IL_{2D-SE} = IL_t^{on} \cdot IL_d^{on}$ on both the bar and cross states. Besides obtaining this symmetry, the 2D-SE forces leakage signals to pass two rings with loss $IL_t^{off} \cdot IL_d^{off}$, resulting in crosstalk levels of 39.7dB:

$$X_{2D-SE}^{cross} = X_{2D-SE}^{bar} = \frac{IL_d^{on} \cdot IL_t^{on}}{IL_d^{off} \cdot IL_t^{off}} = \frac{1}{ER_d \cdot ER_t} \quad (2)$$

III. INTERCONNECTION ARCHITECTURES

We study three silicon photonic interconnections to build an on-chip fabric for board to board interconnections **without considering E-O conversion issues**. Benes networks (in two flavors, respectively using 2B-SEs or 2D-SEs) are chosen because they exhibit, among non-blocking 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. We limit our analysis to single wavelength operation: all TXs use, and all rings resonate at the same wavelength. The analysis can be extended to a WDM scenario, by assuming wavelength striping, with WDM channels fitting exactly the periodical MRs' transfer function [1].




Fig. 3: Interconnection architectures and critical signal paths

A. Scalability of microring-based crossbars

A $N \times N$ crossbar can be implemented as a matrix of N^2 1x2 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 $C_{xbar}(N) = N^2$. More in detail, the worst case path is the one connecting input *in1* to output *outN*, which gives a useful output power $P_{out} = \frac{P_{in}}{IL_{xbar}}$, where $IL_{xbar} = IL_d^{on} \cdot (IL_t^{on})^{2(N-1)}$. For each of the $N-1$ remaining 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 $P_{CT} = \frac{P_{in}}{IL_t^{off} \cdot (IL_t^{on})^{N-2}}$, 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 $P_{xtalk} = \frac{P_{CT}}{IL_d^{off}} \sum_{k=0}^{N-2} \left(\frac{1}{IL_t^{on}} \right)^k$. Finally, normalizing the crosstalk to the useful output power P_{out} , and recalling ER_d and ER_t definitions, we have:

$$X_{xbar} = \frac{(IL_t^{on})^{N-1}}{ER_d \cdot ER_t} \sum_{k=0}^{N-2} \left(\frac{1}{IL_t^{on}} \right)^k \quad (3)$$

B. Scalability of microring-based Benes

Multi-stage $N \times N$ Benes networks show a much lower complexity than crossbars, equal to $C_{Benes}(N) = \frac{N}{2} \cdot C_{SE} \cdot S(N)$ 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 2x2 SE used. Thus, C_{SE} is either equal to $C_{SE} = 2$ or $C_{SE} = 4$ for the 2B-SE or the 2D-SE respectively. The scalability of a Benes network depends on the number of network stages optical signals have to cross from inputs to outputs. When using 2B-SEs, being $IL^{bar} > IL^{cross}$ and $X^{bar} < X^{cross}$, 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 a 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 $P_j = \frac{P_{j-1}}{IL^{bar}}$; thus, after $S(N)$ stages, $P_{S(N)} = \frac{P_{in}}{(IL^{bar})^{S(N)}}$. Conversely, the XWCP occurs when input signals go through cross-state SEs only. At stage j , the output crosstalk power is $X_j = X_{j-1} + X^{cross}$ (considering the worst case in which the powers at the two input ports are equal). After $S(N)$ stages:

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

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

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 an 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 vs. R_b of 13.5 dB/decade and a reference value of -26 dBm at 10 Gbps. 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 can be limited either by losses or crosstalk penalties. Fig. 4, reports crosstalk and loss penalties vs. the total number of ports N . To check the analytical formulation, we performed a TDMS (Time-Domain Monte-Carlo Simulation) to estimate the crosstalk penalty in the Crossbar IX case. The results (see Fig. 4) 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 transmitter bitrate. 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 bitrates (i.e., 1 Gbps) 2D-SE Benes reach the highest aggregated bandwidth at about 8.1 Tbps, 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 Tbps for the crossbar and 1.6 Tbps for the 2D-SE Benes with 100 Gbps 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




Fig. 4: Power penalties for all fabrics




Fig. 5: Aggregate bitrate and complexity comparison

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 Tbps 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. Tech. Letters*, v.20, n.10, May 2008
- [2] Chao Li, A.W.Poon, "Silicon Electro-Optic Switching Based on Coupled-Microring Resonators", *CLEO (Conf. on Lasers and Electro-Optics)*, Baltimore, MD, May 2007
- [3] B.G. Lee, A. Biberman, N. Sherwood-Droz, C.B. Poitras, M. Lipson, K. Bergman, "High-Speed 2x2 Switch for Multi-Wavelength Message Routing in On-Chip Silicon Photonic Networks", *ECOC*, Sept. 2008
- [4] H.Takahashi, K.Oda, H.Toba, "Impact of crosstalk in an arrayed-waveguide multiplexer on N×N optical interconnection", *IEEE J. Light-wave Technology*, v.14, n.6, June 1996, pp.1097-1105
- [5] E.Sackinger, *Broadband Circuits for Optical Fiber Communication*, New York, John Wiley & Sons, 2005
- [6] H.Bukkems, C.Herben, M.Smit, F.Groen, I.Moerman. "Minimization of the loss of intersecting waveguides in InP-based photonic integrated circuits", *IEEE Photon. Tech. Lett.*, v.11, n.11, pp.1420-1422, Nov. 1999
- [7] R. Gaudino, G. A. Gavilanes Castillo, F. Neri, J.M Finochietto, "Simple Optical Fabrics for Scalable Terabit Packet Switches", *IEEE ICC*, Beijing, China, May 2008