

Optical Interconnection Networks Based on Microring Resonators

Original

Optical Interconnection Networks Based on Microring Resonators / Bianco, Andrea; Cuda, Davide; GARRICH ALABARCE, Miquel; Gaudino, Roberto; GAVILANES CASTILLO, GUIDO ALEJANDRO; Giaccone, Paolo; Neri, Fabio. - STAMPA. - (2010), pp. 1-5. (Intervento presentato al convegno IEEE ICC, IEEE International Communications Conference tenutosi a Cape Town, South Africa nel May 2010) [10.1109/ICC.2010.5502490].

Availability:

This version is available at: 11583/2371959 since:

Publisher:

IEEE

Published

DOI:10.1109/ICC.2010.5502490

Terms of use:

This article is made available under terms and conditions as specified in the corresponding bibliographic description in the repository

Publisher copyright

(Article begins on next page)

Optical Interconnection Networks Based on Microring Resonators

A. Bianco*, D. Cuda*, M. Garrich[†], R. Gaudino*, G. Gaviñanes*, P. Giaccone* and F. Neri*

* Dip. di Elettronica, Politecnico di Torino, Italy, Email: {lastname}@polito.it

[†] Universitat Politècnica de Catalunya, Spain, Email: miquel.garrich@estudiant.upc.edu

Abstract—Interconnection networks must transport an always increasing information density and connect a rising number of processing units. Electronic technologies have been able to sustain the traffic growth rate, but are getting close to their physical limits. In this context, optical interconnection networks are becoming progressively more attractive, especially because new photonic devices can be directly integrated in CMOS technology. Indeed, interest in microring resonators as switching components is rising, but their usability in full optical interconnection architectures is still limited by their physical characteristics. Indeed, differently from classical devices used for switching, switching elements based on microring resonators exhibit asymmetric power losses depending on the output ports input signals are directed to. In this paper, we study classical interconnection architectures such as crossbar, Benes and Clos networks exploiting microring resonators as building blocks. Since classical interconnection networks lack either scalability or complexity, we propose two new architectures to improve performance of microring based interconnection networks while keeping a reasonable complexity.

I. INTRODUCTION

The predictions outlined in the “International Roadmap for Semiconductors” [1] show that the most critical performance limitations of on-chip interconnections depend on the latency and power requirements of metal-dielectric wiring, especially for lengths above the millimeter. The failure of copper wires to scale with CMOS transistors has already resulted in a departure from the synchronous single chip/core design paradigm to multi-core or multi-chip systems for both “consumer” and “High Performance Computing” (HPC) systems, which recently approached aggregated bandwidth of Petabit/s. On the one hand, electronic interconnects can not keep up with the continuously increasing bandwidth demand. Indeed, the required parallelism is becoming unsustainable because of electromagnetic compatibility problems, power supply and dissipation requirements. On the other hand, the recent breakthroughs made in CMOS-compatible silicon photonic integration are boosting the penetration of optical technologies in interconnection systems [2], [3]. Photonic technologies can transport a huge information density and their performance are largely independent of the bitrate. Furthermore, they offer the possibility to cover large distances without regeneration.

Among the optical devices that have been recently developed, one of the most promising is the *silicon microring resonator* [4], a small foot-print device suited to a wide range of applications which include signal processing, filtering, delaying or modulating optical signals; thus, microring resonators

have been used also to build sensors, lasers, modulators, switches, memories and slow-light elements. We focus on the use of microring resonators as optical switching elements.

As Fig. 1(a) shows, in microring resonators, an incoming optical signal can be either coupled to the ring (if the input signal wavelength matches the microring’s resonance wavelength) or it can continue along its path (if the input signal wavelength is different from the microring’s resonance wavelength).

Microring resonators show intrinsic asymmetric power penalties, because optical signals coupled to the ring suffer larger power penalties than signals not traversing it. This physical asymmetric behavior implies new challenges in designing optical interconnection networks based on microring resonators, because the signal penalty depends on the paths input signals follows crossing these optical interconnection networks. Indeed, it becomes crucial to minimize the number of times an input signal is coupled to a microring resonator to reduce the overall power penalty while traveling through the interconnection network. On the contrary, in classical interconnection networks (both electronic and photonic) input signals suffer the same power losses independently of the state of the crossed switching element.

In Sec. II, we describe Switching Elements (SE) based on microring resonators that can be used as building blocks to create interconnection networks. In Sec. III, we investigate the scalability and the complexity of crossbar, Benes and Clos networks based on microring SEs; in particular, we propose two new variations of multistage networks to optimize the tradeoff between scalability and complexity. Finally, we present some design evaluation of the proposed solutions in Sec. IV, and draw some conclusions in Sec. V.

II. MICRORING-BASED SWITCHING ELEMENTS

Microring resonators are characterized by a waveguide bent to itself in a circular-like shape, coupled to one or two waveguides or to another microring [4]. Fig. 1(a) shows an example of a microring coupled to two waveguides to build a 1×2 SE; this basic building block is called 1B-SE. Optical signals entering the *input* port can be deflected either to the *drop* port, when the ring is properly tuned to be in resonance with the input signal wavelength, or to the *through* port in the normal non-tuned state. Note that the waveguide intersection introduces negligible scattering losses.

Experimental measurements [4] show that the 1B-SE presents an asymmetric behavior. Indeed, input signals sent

to the drop port experience larger power losses (around 1.4 dB) than signals routed to the through port (around 0.1 dB) because of the propagation inside the ring and the ring-waveguide coupling. Hence, from the power budget point of view, this SE can be in either a High-Loss State (HLS) or in a Low-Loss State (LLS). Furthermore, some power leakage is experienced, being larger on the through port with respect to the drop port (because passing through the ring introduces a power penalty). This leads to an unbalanced coherent crosstalk accumulation when several other SEs are present. In this paper, we assume that the power losses in LLS and that the effect of accumulation of crosstalk due to power leakages are both negligible. These assumptions are reasonable because power penalties of SEs in HLS are much higher than all other power impairments. We don't face the issue of physical modeling of microring resonators; rather, we wish to study interconnection network architectures able to scale to large size by reducing the number of SEs in HLS that optical signals should cross while moving from input to output ports.

Microring resonators can dynamically change their state. Indeed, either exploiting thermal-optic or electro-optic [5] effects or by means of an optical pump [6], it is possible to modulate the microring effective refractive index, hence its resonance frequency. Depending on the technology used, different tuning times as well as different power penalties can be observed. In the remainder of the paper, we assume that microrings are controlled by carrier injection techniques because they ensure a switching time of few hundreds ps [5], suitable to support fast switching. We also assume single wavelength operation: the incoming optical signal is coupled to the ring if its resonance frequency matches the signal frequency. The ring can also be tuned to a different frequency so that the incoming signal is decoupled from the ring.

Building up on the 1B-SE structure, it is possible to design more complex SEs that can be used as building elements in interconnection architectures. Fig. 1(b) depicts a possible implementation of a basic 2×2 SE (called 2B-SE). The 2B-SE is based on two 1B-SEs jointly controlled to provide two switching states: the *bar* state ($\text{in1} \rightarrow \text{out1}$, $\text{in2} \rightarrow \text{out2}$) and the *cross* state ($\text{in1} \rightarrow \text{out2}$, $\text{in2} \rightarrow \text{out1}$). More in detail, in the bar state, each ring deflects the corresponding optical input signal to the drop port of the respective 1B-SE. On the contrary, when the 2B-SE operates in the cross state, each ring lets the corresponding optical input signal pass to the through port of the respective 1B-SE. Hence, also the 2B-SE exhibits an asymmetric behavior, because the bar state is a HLS introducing power losses for both input ports, and the cross state is LLS, with negligible power losses, as experimentally measured in [4].

III. ARCHITECTURES FOR INTERCONNECTION NETWORKS

The SEs presented in Sec. II can be used as building elements to assemble larger interconnection networks. We aim at maximizing interconnects scalability (in terms of port count) and minimizing its complexity C , measured as the number of microrings used to build the interconnects. Due




Fig. 1. Elementary microring-based switching elements

to the asymmetric behavior characterizing SEs, maximizing scalability is equivalent to minimize the maximum number of SEs in HLS (either the drop port for a 1B-SE or the bar-state for a 2B-SE) crossed by optical signals considering every possible input/output connection. We denote by X the maximum number of SEs configured in HLS that an input signal crosses in the optical interconnection network.

In the following we present different solutions to build a microring-based interconnection network with N input/output ports. First, we discuss the properties of the crossbar architecture. The search for non-blocking networks, less complex than the crossbar, naturally leads to multistage interconnection networks. Among the large number of well-known multistage architectures, we then consider in this paper Clos and Benes networks. Finally, to achieve a better trade-off between scalability and complexity, we introduce two variations of multistage networks, combining Benes networks and crossbars: i) the Hybrid Crossbar-Benes (HCB) network and ii) the Hybrid Benes-Crossbar (HBC) network.

A. Crossbar networks based on microrings resonator

Crossbars can be easily built exploiting 1B-SEs: Fig. 2 shows an example of a 4×4 (i.e., $N = 4$) crossbar, in which column waveguides are the input ports and row waveguides are the outputs. The crossbar exhibits the best scalability performance, because each input can be connected to any output crossing a single HLS 1B-SE. Hence, the maximum number of SEs in HLS that any optical signal crosses is $X_{\text{XBAR}} = 1$. However, it is well known that the crossbar requires high complexity, measured in terms of number of microrings needed to create the interconnection network, equal to $C_{\text{XBAR}} = N^2$. As a consequence, the crossbar requires a large footprint and a large number of SEs must be controlled, although via a trivial routing algorithm.

B. Clos networks based on microring resonators

Clos networks are a class of well-known interconnection networks. They are usually employed either when the size (in number of ports) of the interconnection network exceeds the size of the largest feasible crossbar or to reduce the overall network complexity. Indeed, the key advantage of




Fig. 2. 4 × 4 microring-based crossbar




Fig. 3. Basic Clos network

Clos networks is that the number of microrings deployed is much lower than those needed to build a crossbar of the same size. Fig. 3 shows the construction rule to obtain a symmetric three-stage Clos network. The input (output) stage of a $N \times N$ Clos network is composed by k SEs of size $n \times p$ ($p \times n$), with $n = N/k$ denoting the number of ports in each input (output) stage, and p the number of SEs of the middle stage. The number of blocks of the middle stage determines the blocking property of the Clos network. We consider $p = n$, corresponding to the minimum number of middle stages needed to build a rearrangeable (REAR) Clos network [7].

When all SEs are implemented by crossbars, $X_{\text{REAR}} = 3$ and the minimum network cost, obtained for $n = \sqrt{N/2}$, is equal to $C_{\text{REAR}} = 2\sqrt{2}N^{3/2}$.

C. Benes networks based on microring resonators

Benes networks exhibit several advantages, such as a straightforward recursive construction rule, a simple routing and re-arranging algorithm and a complexity asymptotically close to the minimum. Fig. 4 shows an 8×8 Benes network. An $N \times N$ Benes network permits any one-to-one connection between inputs and outputs using a number of stages (columns of 2×2 SEs) equal to $S = 2 \log_2 N - 1$, each stage including $N/2$ 2×2 SEs. Hence, the complexity of a $N \times N$ Benes network scales as $C_{\text{BENES}} = 2N \log_2 N - N$, because each




Fig. 4. 8 × 8 Benes network

2×2 SE includes 2 rings. Even though the Benes network asymptotically exhibits the lowest complexity, $X_{\text{BENES}} = S$ for worst-case paths in a Benes network. Thus, Benes networks lack scalability, because their physical impairments grow with the network depth S .

D. The Hybrid-Crossbar-Benes network

The HCB network is depicted in Fig. 5 and aims at reducing the complexity of a three-stage Clos network substituting the crossbars composing the middle stage of a Clos network with Benes networks. Clearly, from the cost perspective, the best solution is to make the Benes network as large as possible. This condition would lead to make k (the number of ports of the middle stage switching blocks) as large as possible, or, in a rearrangeable network, to have n (the number of ports of the edge stages and the minimum number of middle stage switching blocks) as small as possible. Indeed, for $k = \frac{N}{2}$ the HCB network degenerates to a $N \times N$ Benes network. On the other hand, the optimal solution, scalability wise, is to make the Benes network as small as possible, and this solution degenerates to the Crossbar networks for $k = 1$. From the scalability perspective, an HCB network presents $X_{\text{HCB}} \leq X_{\text{BENES}} + 2 = 2 \log_2 k + 1$. Thus, if we define a target X_T equals to the maximum number of HLS SEs that optical signals are allowed to cross inside the optical interconnection network and substitute $n = \frac{N}{k}$ in X_{HCB} , it is possible to establish the following construction rule:

$$n_{\text{HCB}} \geq \frac{N}{\alpha} \quad (1)$$

with $\alpha = 2^{\frac{X_T - 1}{2}}$.

The value $n_{\text{HCB}} = N/\alpha$ ensures both feasibility (respecting the HLS constraints given by X_T) and minimum complexity. Finally, from Eq. (1) the overall complexity can be expressed as

$$C_{\text{HCB}} = N(2n_{\text{HCB}} - 1) + 2N \log_2 \left(\frac{N}{n_{\text{HCB}}} \right)$$

which ensures a complexity lower than a crossbar for

$$N \geq \alpha(2 - X_T)/(2 - \alpha)$$




Fig. 5. Hybrid Crossbar Benes (HCB) network

E. The Hybrid Benes-Crossbar network

The HBC network is based on the observation that a $N \times N$ Benes network is made of two edge stages of SEs connected to two $(N/2) \times (N/2)$ non-blocking networks in the middle stage. These non-blocking networks are usually Benes networks as well. Indeed, being crossbars non-blocking and optimal in terms of scalability, they can be employed in the middle stage of a multistage network interconnected according to the Benes pattern, using edge stages composed by 2×2 SEs. For instance, in Fig. 6, four 4×4 crossbars are used as middle stages to build a 16×16 network. Crossbars are employed as building blocks for two 8×8 networks (shaded in Fig. 6), which are in turn interconnected to build the 16×16 network.

Let n be the size of crossbars constituting the middle stages to ensure a maximum X_T . Differently from a Benes network, the HBC architecture is always feasible by choosing a proper value for n , ranging from 2 to N . Note that, when $n = 2$ or $n = N$, the HBC architecture degenerates to a Benes or a crossbar network respectively. Finally, also the HBC has $X_T \leq 2 \log_2 \frac{N}{n} + 1$, which leads to the following construction rule:

$$n_{\text{HBC}} \geq \frac{N}{\alpha} \quad (2)$$

with $\alpha = 2^{\frac{X_T-1}{2}}$. The HBC complexity depends on the ratio between N and n , i.e., the smaller the inner crossbars, the lower the HBC complexity. The HBC cost scales as

$$C_{\text{HBC}}(N, n) = 2N \log_2 \left(\frac{N}{n} \right) + Nn$$

From Eq.(2), $n_{\text{HBC}} = N/\alpha$ is the size of the crossbar that ensures minimal complexity for a given target X_T . Finally, the HBC architecture shows a complexity lower than a crossbar for $N \geq (\alpha(1 - X_T)) / (1 - \alpha)$.

IV. COMPLEXITY AND SCALABILITY

In this section we present some design evaluations of the proposed interconnection networks. For the Clos network, the number of SEs in HLS which optical signals should cross does not depend on the blocking condition (which determines the




Fig. 6. Hybrid Benes-Crossbar (HBC) network example




Fig. 7. X as a function of N for crossbar, Clos, Benes and the HBC/HCB networks. For the HCB/HCB, $n = 16$ and $n = 32$ are considered.

number of SEs (p) in the middle stage). Thus, considering non blocking networks, we report results for $n = p$, which is the optimal choice from the cost perspective.

Fig. 7 shows the scalability conditions for the proposed architectures. Due to their construction rule, both crossbar and Clos networks show a constant X equal to 1 and 3, respectively. On the contrary, the Benes network shows the worst scalability performance, because X scales logarithmically with respect to the number of inputs N . The HCB and the HBC solution show the same scalability law, as described by (1) and (2). Results for the feasible configurations $N \geq n$ are presented for the two different cases $n = 16$ and $n = 32$.

Fig. 8 and Fig. 9 show the complexity for the different interconnection networks when varying the number of input/output ports N for two different value of X_T . In these figures we considered $X_T = 7$ and $X_T = 15$ to illustrate two different feasibility conditions which allow to cross a number of HLS SEs which is either quite low or rather high, respectively. The large figure refers to “small” networks (N ranging from 2^4 to 2^{10}), whereas the inset figure refers to “large” networks (N ranging from 2^{10} to 2^{16}). To both satisfy the feasibility condition and to minimize cost, according to (1) and (2), we select $n = N/\alpha$ with $\alpha = 8(128)$ for $X_T = 7(15)$. Note that interconnection networks interconnecting 256 cores have been recently proposed [8], [9] and [10]. Thus, if the current trend persists, even larger interconnection network might be needed.

As expected, complexity is upper- and lower-bounded by crossbars and Benes networks which always show the highest




Fig. 8. Complexity of crossbar, Benes, HBC and HCB networks with $X_T = 7$




Fig. 9. Complexity of crossbar, Benes, HBC and HCB networks with $X_T = 15$

and the lowest cost, respectively. Hybrid architectures are always feasible even for the most demanding target ($X_T = 7$), while for the most relaxed condition ($X_T = 15$) Benes is complexity-convenient until feasible, i.e., up to 256 ports. Indeed, in Fig. 8, for $X_T = 7$, the largest feasible Benes network is 16×16 . Indeed, as the network size increases, the number of HLS SEs that affect the signal is constrained to grow. As a consequence, the feasibility target requires edge (inner) crossbars to be larger for the HCB (HBC) architecture. Finally, Fig. 9 shows that as X_T increases, the size of crossbar SEs used inside the HBC and HCB networks decreases, cutting down the complexity. Note that, HCB and HBC solutions become feasible as the Benes network violates the HLS constraint. The HCB architecture exhibits a larger complexity than HBC network.

V. CONCLUSIONS

In this paper we analyzed scalability and complexity of different interconnection networks based on microring resonators. We initially described basic 1×2 and 2×2 SEs

and we highlighted their asymmetric behavior in terms of power penalties depending on their switching state. Then, we analyzed the effects of these asymmetries on the scalability of crossbar, Benes and Clos networks. To trade scalability and costs, we proposed two architectures, named HBC and HCB, based on different combinations of Benes and crossbar networks.

We showed that HBC and HCB interconnection networks are able to overcome the scalability limitation of Benes networks, presenting, at the same time, a remarkably lower complexity than a crossbar.

It is worth observing that, given the possibility of bounding both impairments and costs, an optical interconnection network based on microring resonators can be conceived following the reasoning described in the paper. Thus, we believe microring resonators can be considered as interesting devices for future photonic interconnection networks.

ACKNOWLEDGMENTS

This work was partially supported by the BONE project, a Network of Excellence funded by the European Commission within the 7th Framework Programme.

REFERENCES

- [1] S. I. Association, "International technology roadmap for semiconductors," http://www.itrs.net/Links/2007ITRS/2007_Chapters/.
- [2] A. V. Krishnamoorthy, X. Z. R. Ho, H. Schwetman, J. Lexau, P. Koka, G. Li, I. Shubin, and J. E. Cunningham, "The integration of silicon photonics and vlsi electronics for computing systems intra-connect," in *Photonic in Switching*, 2009.
- [3] M. Petracca, B. G. Lee, K. Bergman, and L. P. Carloni, "Design exploration of optical interconnection networks for chip multiprocessors," in *Hot Interconnects*, 2008, pp. 31–40.
- [4] B. G. Lee, A. Biberman, N. Sherwood-Droz, C. B. Poitras, M. Lipson, and K. Bergman, "High-speed 2x2 switch for multi-wavelength message routing in on-chip silicon photonic networks," *IEEE/OSA Lightwave Technology, Journal of*, vol. 27, no. 14, pp. 2900–2907, July 2009.
- [5] C. Li and A. Poon, "Silicon electro-optic switching based on coupled-microring resonators," in *Conference on Lasers and Electro-Optics (CLEO)*, May 2007.
- [6] 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 Photonic Technology Letters*, vol. 20, no. 10, pp. 767–769, May 2008.
- [7] A. Pattavina, *Switching Theory, Architectures and Performance in Broadband ATM Networks*. Wiley, 1998.
- [8] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. P. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. G. Beausoleil, and J. H. Ahn, "Corona: System implications of emerging nanophotonic technology," in *International Symposium on Computer Architecture (ISCA 08)*, June 2008.
- [9] Y. Jin, Z. Chang, J. Tang, and W. Hu, "Energy-efficient on-chip photonic hypermesh network for multi-processor communication," in *4th International Workshop on OPS and OCDMA (IWOOC 2009)*, Nov 2009.
- [10] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. Holzwarth, M. Popovic, H. Li, H. Smith, J. Hoyt, F. Krtner, R. Ram, V. Stojanovic, and K. Asanovic, "Building manycore processor-to-dram networks with monolithic silicon photonics," in *Hot Interconnects*, 2008, pp. 21–30.