

Optical interconnection networks based on microring resonators

Original

Optical interconnection networks based on microring resonators / Bianco, Andrea; Cuda, Davide; GARRICH ALABARCE, Miquel; GAVILANES CASTILLO, GUIDO ALEJANDRO; Gaudino, Roberto; Giaccone, Paolo. - In: JOURNAL OF OPTICAL COMMUNICATIONS AND NETWORKING. - ISSN 1943-0620. - STAMPA. - 4:7(2012), pp. 546-556. [10.1364/JOCN.4.000546]

Availability:

This version is available at: 11583/2503364 since: 2016-04-14T11:47:30Z

Publisher:

IEEE

Published

DOI:10.1364/JOCN.4.000546

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

Andrea Bianco^{*}, Davide Cuda[†], Miquel Garrich^{*}, Guido Gavilanes Castillo[‡], Paolo Giaccone^{*}

^{*} Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino, Torino, Italy

[†] Orange Labs, Issy-Les-Moulineaux, France

[‡] Istituto Superiore Mario Boella, Torino, Italy

Abstract—Optical microring resonators can be integrated on a chip to perform switching operations directly in the optical domain. Thus they become a building block to create switching elements in on-chip optical interconnection networks, which promise to overcome some of the limitations of current electronic networks. However, the peculiar asymmetric power losses of microring resonators impose new constraints on the design and control of on-chip optical networks.

In this work, we study the design of multistage interconnection networks optimized for a particular metric that we name *degradation index*, which characterizes the asymmetric behavior of microrings. We also propose a routing control algorithm to maximize the overall throughput, considering the maximum allowed degradation index as a constraint.

Index Terms—Optical interconnections, microring resonators, multistage switching networks, routing algorithm

I. INTRODUCTION

Optical technologies appear as a viable alternative to electronics to build interconnection networks suited for on-chip integration. On the one hand, predictions outlined by the International Technology Roadmap for Semiconductors [1] show that the main performance limitation of feasible on-die electronic interconnections is the length of metal-dielectric wirings, becoming critical for distance above the millimeters. Indeed, the higher the information densities to support, the stricter become the design constraints to be fulfilled, especially in terms of (i) electromagnetic compatibility issues, (ii) maximum distances which electronic signals can cover without regeneration, and (iii) power requirements. On the other hand, recent breakthroughs in CMOS-compatible silicon photonic integration [2] are boosting the penetration of optical technologies into interconnection systems [3], [4]. As a matter of fact, photonic technologies can transport huge information densities, their performance are roughly independent of the bitrate and offer the possibility to cover larger distances without regeneration. Several issues still remain to be solved, both at the component level, e.g., efficient light generation, and at the architectural level, where interconnection architectures exploiting the distinctive features available in the optical domain must be identified. In this paper we focus on the architectural level, studying on-chip optical interconnection networks.

Silicon *microring resonators* represent one of the most promising optical devices to this end. Microring resonators are small foot-print devices, which have already proved to be suited for a wide range of applications, including signal processing, filtering, delaying or modulating optical signals. Furthermore, they have been considered to build sensors, modulators, microlasers and buffers [5], [6].

In this work, we investigate how to design and control interconnection networks based on microring resonators acting as optical Switching Elements (SEs), as proposed in [3], [7]–[9]. In a microring resonator SE (see for example Fig. 1(a)), an incoming optical signal can be either coupled to the ring (if the input signal wavelength is equal to the resonance wavelength of the microring) or it can continue along its path (if the input signal wavelength is different from the resonance wavelength). This allows to switch an incoming signal to a different output port, depending on the ring resonance wavelength.

In classical interconnection networks, input signals suffer power losses independent of the state of the SEs in the interconnection network. On the contrary, SEs based on microring resonators show an intrinsic asymmetric state-dependent behavior because optical signals coupled/not-coupled to the ring suffer different power penalties, depending on the specific microring design [9]. To enhance architecture scalability, it is crucial to minimize the power degradation experienced by *any* signal traveling through the interconnection network. To this aim, we propose different architecture designs and suitable algorithms to control the SEs configuration.

Our major contributions are: (i) two multistage interconnection networks that minimize the power losses, (ii) a mirroring technique to scale up the size of the network, and (iii) a control algorithm that considers the specific physical impairments. All these contributions allow to design and control the whole switching networks, leading to good performance in terms of complexity and port scalability. Preliminary results were presented in [10] and [11].

The remainder of the paper is organized as follows. In Sec. II, we first describe simple microring-based SEs that can be used as building blocks to create larger interconnection networks. In Sec. III, we define the *degradation index* and we investigate the scalability properties of three classical multistage networks based on microring SEs in terms of area (measured as the number of deployed microrings) and degradation index. To optimize the tradeoff between area and degradation index, we also propose two multistage networks in Sec. IV and the mirroring technique in Sec. V. In Sec. VI we present a variation of the classical Paull’s algorithm for Clos networks [12], to configure microring-based switching fabrics and drastically reduce the degradation index. Finally, we present some design evaluation of the proposed solutions in Sec. VII, and we draw some conclusions in Sec. VIII.

II. MICRORING-BASED SWITCHING ELEMENTS

Microring resonators are basically composed by a waveguide bent to itself in a circular shape, coupled to one or

two waveguides or to another microring. 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 either deflected to the *drop* port, when the ring is properly tuned (or in resonance) with the input signal wavelength, or continue along the input waveguide towards the *through* port in untuned state.

The resonance frequency can be modified exploiting either thermal-optical effects [13], carrier-injection [14], [15] or optical pumps [16]. Depending on the used technology, 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 (short) switching times (few hundred ps according to [14]). We also assume a single wavelength operation; the incoming optical signal is coupled (or not) to the ring, depending on the wavelength to which the ring is tuned. All the proposed architectures can be extended to a Wavelength Division Multiplexing (WDM) scenario, but this is outside the scope of this paper.

Dealing with a physical model of microrings (see for example [8]) is outside the scope of this work. Microring based 1B-SEs present an asymmetric power penalty depending on the switching state. According to [9], the design of the microring resonator is based on a tradeoff between different figures of merit. More precisely, a large coupling coefficient between microring and adjacent waveguides leads to low power degradation to signals deflected to the drop port but high degradation to signals sent towards the through port. On the contrary, a small coupling coefficient between microring and adjacent waveguides leads to high degradation to signals deflected to the drop port but low degradation to signals sent towards the through port. For instance, in the microrings controlled through laser pumps [16], input signals sent to the drop port experience larger power losses (around 1.5 dB) than signals routed to the through port (with negligible losses) because of the small coupling coefficient between microring and drop/through waveguides. Therefore, we describe the loss behavior of the 1B-SE with a simplified model based on two loss states, assuming either a High-Loss State (HLS), or a Low-Loss State (LLS).

Extending the 1B-SE structure, it is possible to design more complex SEs that can be used as elementary switching elements in large interconnection networks. Fig. 1(b) depicts a possible implementation of a basic 2×2 SE (called 2B-SE). The 2B-SE exploits two 1B-SEs jointly controlled to provide two switching states. In the *bar* state ($in_0 \rightarrow out_0$, $in_1 \rightarrow out_1$), each ring deflects the corresponding optical input signal to the drop port of the respective 1B-SE. In the *cross* state ($in_0 \rightarrow out_1$, $in_1 \rightarrow out_0$), each ring lets the corresponding optical input signal pass to the through port of the respective 1B-SE. Also the 2B-SE can be modeled as a SE with two opposite loss states. Without loss of generality, we assume the bar state as a HLS, and the cross state as a LLS, coherently with either the experimental measurements of [7], or the results of [9] for a small coupling coefficient between microring and adjacent waveguides. Note that the penalty due to the waveguide crossing can be minimized applying an expansion technique to achieve a 0.2 dB of loss [17]. In [18], the loss due to the waveguide crossing is further reduced to 0.16 dB applying a double-etched technique.

We propose a modified version of the 2B-SE, useful to optimize the design of multistage interconnection networks.

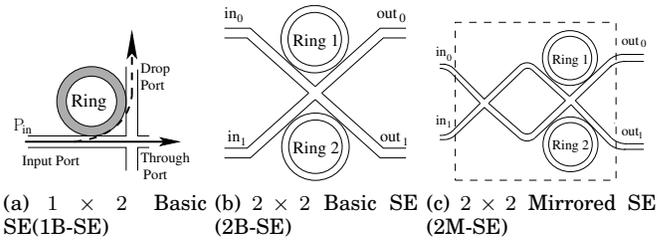


Fig. 1. Elementary microring-based Switching Elements (SE)

Fig. 1(c) shows a 2×2 Mirrored-SE (called 2M-SE). By cross-connecting the input ports, the bar state is now realized by setting up the internal 2×2 SE in the cross state (LLS), whereas the cross state is achieved with an internal bar state (HLS), thus exchanging the LLS with the HLS, with respect to a 2B-SE. The layout of the 2M-SE introduces additional bendings compared with the 2B-SE. However, the corresponding loss can be considered negligible for our analysis, because its value is approximately 0.005 dB for a 90° bend with a $6.5 \mu\text{m}$ radius [6].

Note that our characterization of the SE asymmetric loss behavior can be applied to other optical switching technologies. As an example, the Mach-Zehnder (MZ) based optical 2×2 switch presented in [19] shows a bar-state transmission loss 1 dB higher than the cross-state transmission loss. Similarly, [20] describes a MZ-based optical 2×2 switch with a bar-state transmission loss 0.7 dB higher than the cross-state transmission loss.

In this work we neglect the signal degradation due to crosstalk. Note that [9] shows that the effect of crosstalk depends on the specific switching configuration and on the coupling coefficients between the optical guides. Similarly to our assumptions, the signal impairment is asymmetric for each SE state, although it depends also on the traffic sent through the fabric. Following the approach described in [8], loss and crosstalk penalties could be integrated into a single “impairment index”, thus permitting to extend the proposed methodology to devise a crosstalk-aware design. We leave this extension for future work. Instead, we focus on the design of interconnection network architectures able to scale to large size by reducing the number of SEs in HLS that optical signals should cross while traveling from input to output ports. Our methodology is independent of the specific switching configuration which creates either one of the loss states (HLS or LLS).

III. BASIC MULTISTAGE ARCHITECTURES

The SEs presented in Sec. II is used as building elements to assemble a larger interconnection network with N input and N output ports. Even if we borrow some terminology from the circuit-switching domain, the reference scenario is packet switching. The interconnection network must be configured to support a given set of input-output pairs defined by a permutation π : input i must be connected to output $\pi(i)$. The total number of possible complete permutations is $N!$. When we say that a connection from i to $\pi(i)$ must be established, we mean that a data packet from input i must be transferred to output $\pi(i)$ according to the decision of a packet scheduler. The sequence of SEs traversed by the connection defines a path in the switching network, computed by a proper routing algorithm.

We aim at maximizing the scalability of the architecture for large N by considering both the area and the power losses. The area, denoted by C , is evaluated as the *total number of microrings* present in the interconnection network.

The power losses are evaluated through the maximum number of SEs configured in HLS that the input signal must cross in the optical interconnection, considering all possible input-output permutations, i.e. taking a worst-case approach. Such number is denoted by X and referred to as the *degradation index*. The architecture with the best scalability is achieved by minimizing C and X for a given N . More precisely, we consider a network design constrained by a given maximum degradation index \hat{X} , equal to the maximum number of HLS SEs that optical signals are allowed to cross. Whenever $X \leq \hat{X}$, then all the possible input-output permutations can be established and the corresponding architecture is said to be feasible. On the contrary, when $X > \hat{X}$, there exist some permutations for which some input-output connections cannot be established, because the corresponding path would violate the target \hat{X} . In such a case, the connection is said to be blocked, and the corresponding packet cannot be sent across the switching fabric.

In the following we will consider three basic architectures for interconnection networks: crossbars, Clos networks and Benes networks. Then, we will propose some hybrid architectures, to combine the different tradeoff between area and degradation index offered by each basic architecture.

A. Crossbar networks based on microring resonators

Firstly, we discuss the properties of the crossbar (XBAR) architecture, which will be used also as basic building block for the architectures that will be later proposed. Crossbars can be built exploiting 1B-SEs: Fig. 2 shows an example of crossbar, in which column waveguides are connected to the input ports and row waveguides to the outputs. The crossbar exhibits the best scalability in terms of degradation index, because each input can be connected to any output crossing a single 1B-SE in HLS. Hence, the maximum number of SEs in HLS that any optical signal crosses is $X_{\text{XBAR}}(N) = 1$. Routing is trivial. However, the crossbar exhibits the worst scalability in terms of area, because $C_{\text{XBAR}}(N) = N^2$.

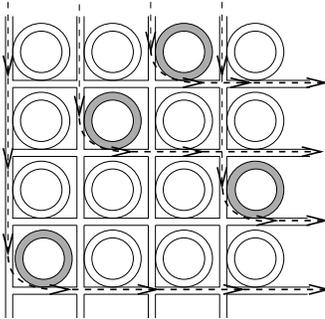


Fig. 2. 4×4 microring-based crossbar connecting $1 \rightarrow 4$, $2 \rightarrow 2$, $3 \rightarrow 1$ and $4 \rightarrow 3$

B. Clos networks based on microring resonators

The definition of networks with lower number of SEs than the crossbar, naturally leads to multistage interconnection

networks, like the well known Clos networks [12]. Fig. 3 shows a symmetric three-stage Clos network: each switching module (SM) of the first stage and of the third stage is connected with all the SMs of the middle stage. The first design parameter is the number of SMs in the first and the third stage, denoted by k . The second design parameter is the number of SMs in the second stage, which affects the blocking property. In the following, we consider $n = N/k$ SMs in the middle stage, because this is the minimum n that guarantees a non-blocking rearrangeable network¹. As a consequence, the first and third stage SMs are of size $n \times n$, and the second stage SMs are of size $k \times k$. When all SMs

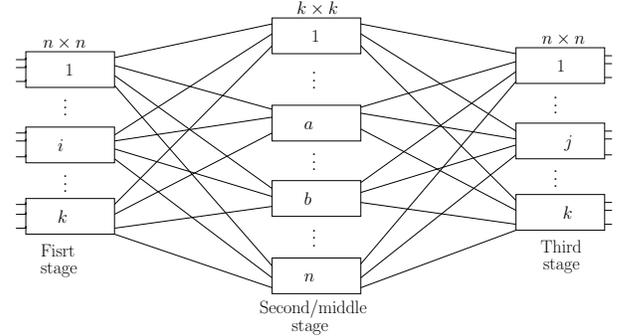


Fig. 3. Rearrangeable non-blocking Clos network

are implemented by crossbars, $X_{\text{CLOS}} = 3$ and the minimum network area, obtained for $n = \sqrt{N/2}$ as shown in [12], is equal to $C_{\text{CLOS}}(N) = 2\sqrt{2}N^{3/2}$.

C. Benes networks based on microring resonators

Benes networks are Clos networks that are recursively constructed with basic SMs of size 2×2 ; thus $N = 2^h$ for some integer h . From the area point of view, they offer the best scalability, being asymptotically optimal [12]. Indeed, an $N \times N$ Benes network has a number of stages (columns of 2B-SEs) equal to $2 \log_2 N - 1$, each stage including $N/2$ basic 2B-SEs. Hence, the area scales as

$$C_{\text{BENES}}(N) = 2N \log_2 N - N \quad (1)$$

because each 2×2 SE includes 2 rings. In terms of degradation index,

$$X_{\text{BENES}}(N) = 2 \log_2 N - 1 \quad (2)$$

since, in the worst case, there exists a path that passes through a HLS SE in each stage. Hence, Benes networks show a poor scalability in terms of degradation index, since X_{BENES} grows with the size N ; given a maximum value \hat{X} for the degradation index, it is impossible to build networks with size $N > 2^{(\hat{X}+1)/2}$.

IV. HYBRID MULTISTAGE ARCHITECTURES

The upper part of Table I provides a synoptic overview of the area and the degradation index for the basic architectures. To improve scalability for large N , we now combine

¹As a reminder, a switching network is defined non-blocking if it is always possible to add a new connection from any free input to any free output port. A non-blocking network is defined rearrangeable, if it is possible to add such new connection by reconfiguring the pre-existing paths due to the previous connections [12]

TABLE I
COMPARISON BETWEEN DIFFERENT NETWORK ARCHITECTURES,
HAVING DEFINED $\hat{k} = 2^{(\hat{X}-1)/2}$.

Network	Area	Degradation index
Crossbar	N^2	$1 \leq \hat{X}$
Clos	$2\sqrt{2}N^{\frac{3}{2}}$	$3 \leq \hat{X}$
Benes	$2N \log_2 N - N$	$2 \log_2 N - 1 \leq \hat{X}$
M-Benes	$4N \log_2 N$	$\log_2 N \leq \hat{X}$
HCB	$2N^2/\hat{k} + N(\hat{X} - 2)$	$3 \leq \hat{X} \leq 2 \log_2 N - 1$
M-HCB	$4N^2/\hat{k}^2 + 2N(2\hat{X} - 3)$	$3 \leq \hat{X} \leq \log_2 N$
HBC	$N^2/\hat{k} + N(\hat{X} - 1)$	$1 \leq \hat{X} \leq 2 \log_2 N - 1$
M-HBC	$4N^2/\hat{k}^2 + 2N(\hat{X} - 2)$	$3 \leq \hat{X} \leq \log_2 N + 1$

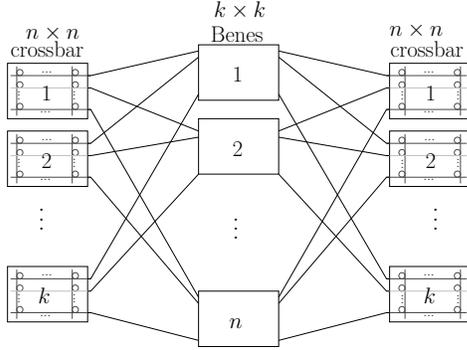


Fig. 4. Hybrid Crossbar-Benes (HCB) network.

the Clos network architecture (based on simple crossbar with low degradation index) with the Benes network (with minimum area), considering two possible variants: *i*) the Hybrid Crossbar-Benes network and *ii*) the Hybrid Benes-Crossbar network.

A. The Hybrid Crossbar-Benes (HCB) architecture

The HCB network is depicted in Fig. 4 and consists of a Clos network in which the middle-stage modules are implemented through Benes networks, instead of crossbars. Like the original Clos network, the HCB network is constrained to $N = nk$, being $n \geq 2$ the size of the crossbars, and $k = 2^h$ the size of the Benes networks for some integer $h \geq 1$. Clearly, from the area perspective, the best solution is to make the Benes network as large as possible, up to $k = N/2$, when the HCB network degenerates into a Benes network. On the contrary, from the degradation index perspective, since

$$X_{\text{HCB}} = 2 \log_2(N/n) + 1 \quad (3)$$

we must increase the crossbar size n . Thus, given a target $\hat{X} \geq 3$, it must be $X_{\text{HCB}} \leq \hat{X}$ and $n \geq N/\hat{k}$, having defined $\hat{k} = 2^{(\hat{X}-1)/2}$. The value $n = N/\hat{k}$ ensures both feasibility and minimum area. Finally,

$$\begin{aligned} C_{\text{HCB}}(N, \hat{X}) &= 2\hat{k}C_{\text{XBAR}}(n) + nC_{\text{BENES}}(\hat{k}) = \\ &= 2\frac{N^2}{\hat{k}} + N(2 \log_2 \hat{k} - 1) = 2\frac{N^2}{\hat{k}} + N(\hat{X} - 2) = \\ &= \frac{N^2}{2^{(\hat{X}-3)/2}} + N(\hat{X} - 2) \quad (4) \end{aligned}$$

for $3 \leq \hat{X} \leq 2 \log_2 N - 1$. Note that (4) is lower than a crossbar for $N \geq (\hat{X} - 2)/(1 - 2^{-(\hat{X}-3)/2})$.

B. The Hybrid Benes-Crossbar (HBC) architecture

The HBC network is depicted in Fig. 5 (with $N = n2^h$) and consists of a Benes network factorized until a certain level $h \geq 0$, when the middle stage is substituted by $k = 2^h$ crossbars of size n .

Again, from the area perspective, we must increase the ‘‘Benes part’’ up to $k = N/2$ when the HBC degenerates into a Benes network. However, from the degradation index point of view, since

$$X_{\text{HBC}} = 2h + 1 = 2 \log_2(N/n) + 1 \quad (5)$$

the best solution is to decrease the level of factorization h (up to $h = 0$ when the HBC degenerates into a crossbar). Thus, to satisfy a given target $\hat{X} \geq 1$, it must be $X_{\text{HBC}} \leq \hat{X}$ and $n \geq N/\hat{k}$, having defined $\hat{k} = 2^{\hat{h}}$ and $\hat{h} = (\hat{X} - 1)/2$. The value $n = N/\hat{k}$ ensures both feasibility and minimum area. Now the HBC area scales as:

$$\begin{aligned} C_{\text{HBC}}(N, \hat{X}) &= \hat{k}C_{\text{XBAR}}(n) + 2(2\hat{h})(N/2) = \\ &= \frac{N^2}{\hat{k}} + N(\hat{X} - 1) = \frac{N^2}{2^{(\hat{X}-1)/2}} + N(\hat{X} - 1) \quad (6) \end{aligned}$$

for $1 \leq \hat{X} \leq 2 \log_2 N - 1$. Note that (6) is asymptotically half the area of the HCB and it is lower than a crossbar for $N \geq (\hat{X} - 1)/(1 - 2^{-(\hat{X}-1)/2})$.

V. MIRRORED ARCHITECTURES

To further reduce the degradation index X or to achieve larger N for a given maximum degradation index \hat{X} , we propose the mirroring technique that exploits the spatial dimension, as shown in Fig. 6(a). We use two different switching planes, topologically identical: the *normal* plane is built with 2B-SEs only, whereas the *mirrored* plane with 2M-SEs only.

Depending on the architecture of the plane selector, the inputs are connected to both planes by means of either one of three solutions: *i*) the 1B-SE (Fig. 1(a)) which implicitly introduces a different degradation index to the signal depending on the plane, or *ii*) the plane selector depicted in Fig. 6(b), which introduces a constant degradation index to both planes, or *iii*), when possible, *integrating* the selector with an extension of the first stage of the architecture. For large N , in the *i*) and *ii*) plane selectors, the number of waveguide crossings may be significant, due to the connection from/to input/output ports to/from the fabric (see the solid and dashed lines in the left hand side of Fig. 6(a)). To

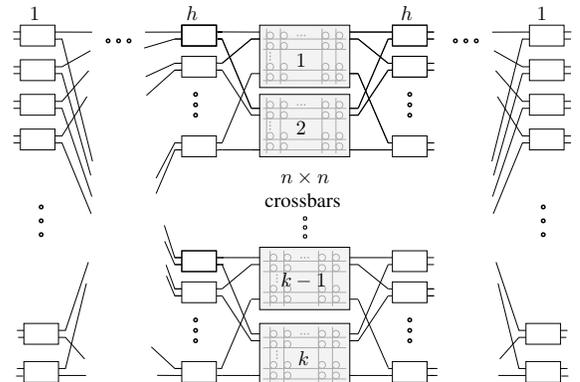


Fig. 5. Hybrid Benes-Crossbar (HBC) network.

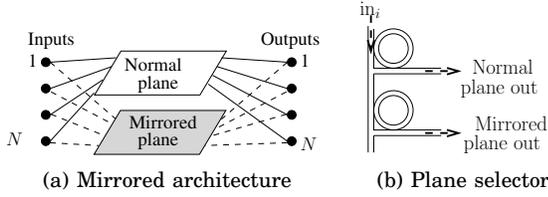


Fig. 6. Mirrored architecture and plane selector

address this issue we propose to adopt a vertically coupled microring resonator [21] for the 1B-SE in i) and for the microring connected to the normal plane in ii). By doing so, input ports are connected to the normal plane through waveguides on a “lower busline level” (solid lines in Fig.6(a)), and to the mirrored plane through waveguides in an “upper busline level” (dashed lines in Fig. 6(a)). According to [21], the required fabrication process for the two busline levels is feasible and would avoid power penalties due to waveguide crossing in our architecture. Since a similar design can be exploited with the couplers from the fabric to the outputs, we neglect losses due to waveguide crossing in our model.

We now consider the two fabrics, one for each plane. Since a HLS of a SE in one plane corresponds to a LLS in the other plane, a generic path that crosses X 2B-SEs in HLS in the normal plane, will cross $(S - X)$ 2M-SEs in HLS in the mirrored plane, where S is the total number of traversed 2×2 SEs in each plane. This observation leads to a simple routing algorithm. After computing a path for the normal plane, whenever $X > S/2$, the routing algorithm will choose the path along the mirrored plane, i.e. the one with lowest degradation index; otherwise the path will be routed across the normal plane. As a consequence, we can start from a plane characterized by a degradation index X and build the whole network with degradation index:

$$X_M = \lfloor X/2 \rfloor + X_S \quad (7)$$

where X_S is the degradation index introduced by the selector; as a consequence, the degradation index is roughly halved by doubling the area of the “Benes part” of the network, constituted only by 2×2 SEs (i.e., doubling the number of 2×2 SEs). This allows to scale the size of the network given the same \hat{X} , as shown below for each architecture.

From the routing point of view, computing the path in the mirrored architecture has the same complexity as computing it for just a single plane, because of the identical topology of the normal plane and the mirrored one.

Note that mirroring the crossbars (based on 1B-SEs) does not provide any advantage in terms of degradation index. Two possible solutions to build the crossbars can be envisaged: i) the crossbars are replicated in both planes; ii) each crossbar is shared between the two planes, by adding one plane selector at the inputs of each crossbar, and one passive coupler at the outputs. Solution i) has the advantage of avoiding additional degradation index, but it introduces the area cost of replicating the crossbars. On the contrary, solution ii) shows a larger degradation index, but this is compensated by a smaller area. In the following, we considered solution i), whereas solution ii) is left for future works.

The mirroring technique will be denoted with the prefix *M*- and it is applied to Benes, HCB and HBC networks.

A. Mirrored Benes networks

In the M-Benes network, the selector is either a 1B-SE or the plane selector of Fig. 6(b). On the one hand, the 1B-SE introduces an asymmetric behavior in terms of degradation index, and makes more complex the layout to connect directly the drop port and the through port to each of the two planes. On the other hand, the plane selector adds a constant degradation index of one extra 1B-SE in HLS to all the paths, but it is simpler to implement due to its layout suited for an homogeneous selection between both planes. Therefore, we will assume always the plane selector for the M-Benes.

Since the Benes network presents an odd number of stages, when we build a plane with a maximum degradation index X_{BENES} , from (7) we obtain a M-Benes network with:

$$X_{\text{M-BENES}} = \frac{X_{\text{BENES}} - 1}{2} + 1 = \log_2 N \quad (8)$$

Let N be the maximum size of a Benes network satisfying \hat{X} ; by (2), $\hat{X} = 2 \log_2 N - 1$. A mirrored version satisfying \hat{X} is built with two Benes networks in which the maximum degradation index is relaxed up to $2\hat{X} - 1$, according to (8). Hence, such Benes networks can have up to N' ports compatible with $2\hat{X} - 1 = 2 \log_2 N' - 1$. By simple algebra, it can be shown that $N' = N^2/2$, i.e. the mirroring technique allows to scale the network by a quadratic factor.

The final area is simply obtained by considering the plane selectors and the two Benes networks:

$$C_{\text{M-BENES}}(N) = 2C_{\text{BENES}}(N) + 2N = 4N \log_2 N$$

B. Mirrored HBC networks

The M-HBC network requires a selector due to the Benes part at the edges of the network. Similarly to the M-Benes network, we assume the plane selector of Fig. 6(b).

Let X_{HBC} be the maximum degradation index in a single plane HBC network. Due to the even number of Benes stages that compose the HBC network and the single crossbar stage in the middle, the degradation index for the M-HBC network is:

$$X_{\text{M-HBC}} = \frac{X_{\text{HBC}} + 1}{2} + 1 = \log_2 \frac{N}{n} + 2 \quad (9)$$

Let \hat{X} be the maximum degradation index satisfied in a HBC network. According to (9), we can build a M-HBC in which the maximum degradation index is relaxed up to $2\hat{X} - 3$. Therefore, recalling that $\hat{k} = 2^{(\hat{X}-1)/2}$, the area scales as:

$$C_{\text{M-HBC}}(N, \hat{X}) = 2C_{\text{HBC}}(N, 2\hat{X} - 3) + 2N = 4 \frac{N^2}{\hat{k}^2} + 2N(\hat{X} - 2) = \frac{N^2}{2^{\hat{X}-3}} + 2N(\hat{X} - 2) \quad (10)$$

for $3 \leq \hat{X} \leq \log_2 N + 1$. Note that for high values of N , (10) is lower than (6) for a degradation index $\hat{X} > 5$.

C. Mirrored HCB networks

Differently from the mirrored versions of Benes and HBC networks, it is possible to *integrate* the plane selector in the first stage crossbar. Indeed, in a mirrored HCB network, the first stage is composed by crossbars of size $n \times n$. The plane selector is connected to two crossbars (one for each plane) and allows to connect each input port with $2n$ output ports (n for each plane). This observation suggests a possible way to integrate n plane selectors and two $n \times n$ first-stage crossbars

with just a single $n \times (2n)$ crossbar, reducing the area by $2N$ and the degradation index by one.

Let X_{HCB} be the maximum degradation index in a single plane HCB network. Due to the odd number of Benes stages that compose the HCB network and the crossbars stages at the edges, we obtain the following degradation index for the M-HCB:

$$X_{\text{M-HCB}} = \frac{X_{\text{HCB}} + 1}{2} = \log_2 \frac{N}{n} + 1 \quad (11)$$

Let \hat{X} be the maximum degradation index satisfied in a HCB network. According to (11), we can build a M-HCB in which the maximum degradation index is relaxed up to $2\hat{X} - 1$. Therefore, recalling that $\hat{k} = 2^{(\hat{X}-1)/2}$, the area scales as:

$$C_{\text{M-HCB}}(N, \hat{X}) = 2C_{\text{HCB}}(N, 2\hat{X} - 1) = 4\frac{N^2}{\hat{k}^2} + 2N(2\hat{X} - 3) = \frac{N^2}{2^{\hat{X}-3}} + 2N(2\hat{X} - 3) \quad (12)$$

for $3 \leq \hat{X} \leq \log_2 N + 1$. Note that for high values of N , (12) is lower than (4) for a degradation index $\hat{X} > 3$.

VI. POWER-PENALTY-AWARE ROUTING ALGORITHM

In previous Sections III-V we have proposed different network architectures, whose area and degradation index have been summarized in Table I. Independently of the configuration algorithm, both the crossbars and the Clos networks experience a fixed degradation index: $X_{\text{XBAR}} = 1$ and $X_{\text{CLOS}} = 3$. On the other hand, for all the other networks considered in this paper, the degradation index depends on the path chosen by the routing algorithm to establish the required connections. In this section, we show the design of a routing algorithm, aware of HLS and LLS states, that reduces the degradation index.

Aim of the routing algorithm is to configure the state of each SE to satisfy any given input-output permutation π , as defined in Sec. III. We consider the well-known Paull algorithm [12] which has been designed to configure a three-stage Clos network, but it is also suitable for multi-stage interconnection networks based on recursive Clos construction. At each recursion level, the network is abstracted as an equivalent three stage Clos network.

Referring to Fig. 3, consider a basic Clos network with I-stage and III-stage switching modules (SMs) of size $n \times n$; as a consequence, n modules are present in the II stage. The algorithm starts from the $N \times N$ network and from a given permutation π of size N , and then it computes the configuration of all the $2k + n$ SMs to establish the paths that route all the connections in π . Then, if the middle-stage SMs are, internally, Clos networks, the Paul algorithm is applied recursively to each individual $k \times k$ SM to establish all the required k connections.

Consider now just a basic Clos network that must be configured according to π . The Paull algorithm works in an incremental way, considering the inputs in an arbitrary order. When input x is considered, the algorithm computes the path to connect x to output $\pi(x)$ by configuring (see Fig. 3): (i) the I-stage SM where input x is located (we denote this as “module i ”), (ii) the III-stage SM where output $\pi(x)$ is located (we denote this as “module j ”), (iii) one or at most two SMs present in the II-stage. By construction, exactly one of the two cases can occur:

- 1) there exists a II-stage module a that can be connected to both modules i and j ; in this case, II-stage module a

is configured to support the connection from its internal input i to its internal output j ; I-stage module i is configured to connect input x to its internal output a ; III-stage module j is configured to connect its internal input a to $\pi(x)$;

- 2) otherwise, there exist two II-stage modules a and b , such that a can be connected to module i , and b can be connected to module j ; in this case, the algorithm moves a set of pre-existing connections from a to b , and another set from b to a , and accordingly recomputes the configurations of the I-stage modules and III-stage modules; then either a or b is used to support the new connection from x to $\pi(x)$, similarly to the previous case.

The Paull algorithm exploits many degrees of freedom to establish the connections given by π : (i) the sequence of inputs considered by adding the connections in π , (ii) the choice of a and b , (iii) the choice of the paths to be moved from/to a and b . Of course, each routing choice affects the degradation index experienced by the paths across the switching fabric.

We propose to modify the Paull algorithm to exploit such degrees of freedom in choosing a and b , among all the n possible II-stage modules, to minimize the degradation index. This modified version of the algorithm will be denoted by *PPA-Paull* (Power-Penalty-Aware Paull algorithm), in contrast with the classical version denoted simply by *Paull*. In the case $n > 2$, the I-stage and III-stage modules are crossbars and the degradation index introduced by them is always two, independently from the choice of the II-stage module: PPA-Paull chooses randomly a II-stage module that is currently available. In the case $n = 2$ (i.e. 2B-SEs at the I-stage and III-stage), the degradation index depends on the state of I-stage module i and III-stage module j . Let us assume that all the inputs and the outputs of the whole Clos network are numbered in increasing order starting from 1 and that the II-stage module a is the upper module, whereas b is the lower one. When an odd input is connected to an odd output, PPA-Paull will choose (if available) b to configure both 2B-SEs at the edges in LLS. Analogously, when an even input is connected to an even output, a will be chosen. Otherwise, either a or b will be chosen at random, since exactly one 2B-SE in the I-stage or III-stage will be in HLS (i.e., the degradation index will be always one).

VII. NUMERICAL EVALUATION

In Sec. VII-A we investigate the scalability of different interconnection networks. In Sec. VII-B we investigate by means of simulations the degradation index reduction offered by the PPA-Paull algorithm with respect to the classical Paull algorithm.

A. Interconnection network design

We compare $N \times N$ interconnection networks by evaluating both the degradation index X and the area (in terms of number of microrings).

Fig. 7 shows the degradation index of different networks in function of N . Crossbars and Clos networks show a fixed degradation index, as discussed in Sections III-A and III-B. On the contrary, the Benes network shows the worst scalability in terms of degradation index, because X scales logarithmically with respect to the number of inputs, coherently with

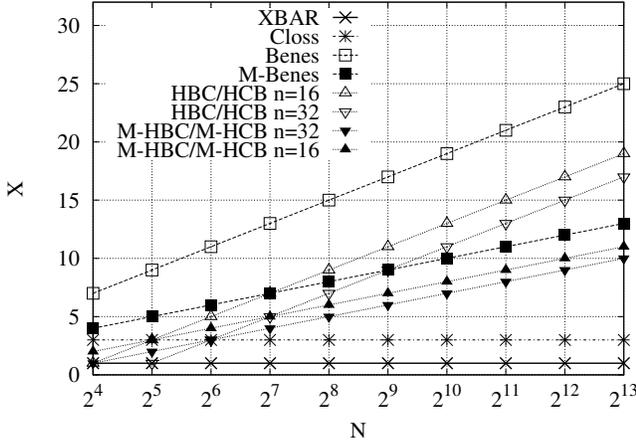
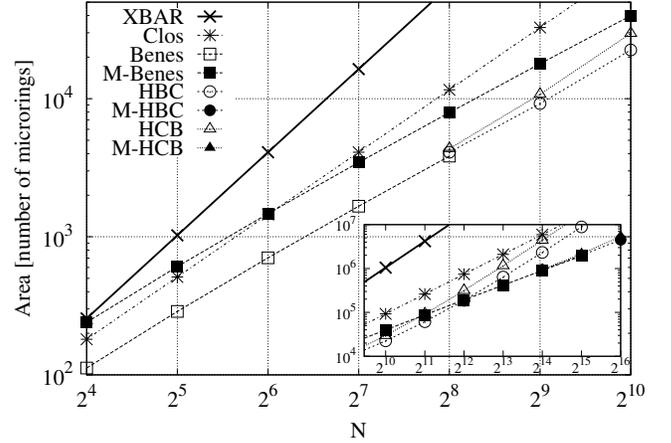
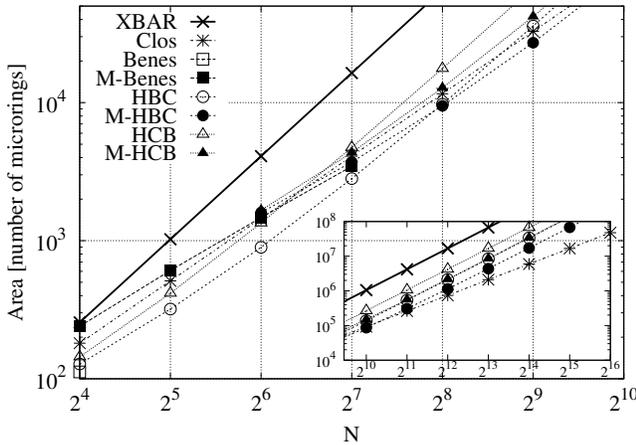


Fig. 7. Degradation index

Fig. 9. Area for maximum degradation index $\hat{X} = 15$ Fig. 8. Area for maximum degradation index $\hat{X} = 7$

(2). The HCB and the HBC networks (for the two different cases $n = 16$ and $n = 32$) show the same scalability law, as described by (3) and (5). As shown in Sec. V, the mirrored technique roughly reduces X by a factor of two. Thus, if we impose a maximum degradation index \hat{X} , it is possible to build Benes networks with a number of ports equal to $N^2/2$ instead of N ; similar gains apply to the Benes part of the other networks. Note that this advantage appears only for large N , as predicted by our models.

Fig. 8 and Fig. 9 show the area for different interconnection networks for two different values of maximum power penalty; the points refer only to feasible configurations, i.e. compatible with the given target \hat{X} . The main figure refers to smaller networks (N ranging from 2^4 to 2^{10}), whereas the inset figure refers to larger networks (N from 2^{10} to 2^{16}).

In general, the crossbar always shows the highest area, while Benes networks always the lowest one, whenever feasible.

Fig. 8 shows that the only feasible Benes network, when $\hat{X} = 7$, is for $N = 16$. Instead, exploiting the mirroring technique, networks up to $N = 128$ can be built. Clos networks show the lowest area for very large N . Indeed, as the network size increases, the worst case degradation index grows. As a consequence, when \hat{X} is smaller, edge or inner crossbars in HBC/HCB networks must be larger. On

the contrary, Fig. 9 shows that, if the target \hat{X} is larger, the size of crossbars inside the HBC and HCB networks decreases, reducing the overall area. Mirrored hybrid architectures reduce by a square factor the size of the internal crossbars. For low values of \hat{X} , the mirroring technique becomes area-effective for smaller N . HCB and HBC networks are feasible whenever the Benes network violates the HLS constraint. Among the two hybrid solutions, the HCB architecture exhibits a larger area than HBC network, and similar observations hold for their mirrored solutions.

B. The performance of the routing algorithm

We compare the degradation index of the PPA-Paull algorithm with respect to the classical version of the Paull algorithm. We report results only for Benes networks, which provide an upper bound on the degradation index experienced by the other architectures.

We assume a microring based interconnection network, with synchronous operations, i.e., time is divided into intervals of fixed duration (timeslots) and the network transfers data units of fixed size (denoted as “cells”); the timeslot duration is equal to the transmission time of a cell. In the case of variable-size packets, incoming packets are chopped into cells, while outputs reassemble all the cells belonging to the same packet. We consider a uniform traffic scenario, with ρ being the average load at each input port. At each timeslot, we generate an input-output permutation π in which each input is active with probability ρ . Starting from a random input and considering all the other inputs in a sequential fashion, a sequence of consecutive connections is generated according to the rule: if input i is active, it is connected to $\pi(i)$. The routing algorithm adds each connection at the time in an incremental way, during the same timeslot.

Each path computed by the algorithm will result in a certain power penalty index; if such index is larger than \hat{X} , the corresponding connection is blocked. To compare the routing algorithms, we measure the blocking probability $P_B(\hat{X})$ as the average fraction of blocked input-output pairs over the number of active inputs. As a complementary measurement, the throughput is evaluated as the average number of connections that are established without blocking in a generic timeslot; the maximum throughput is equal to the average load ρ and it is reached when a connection

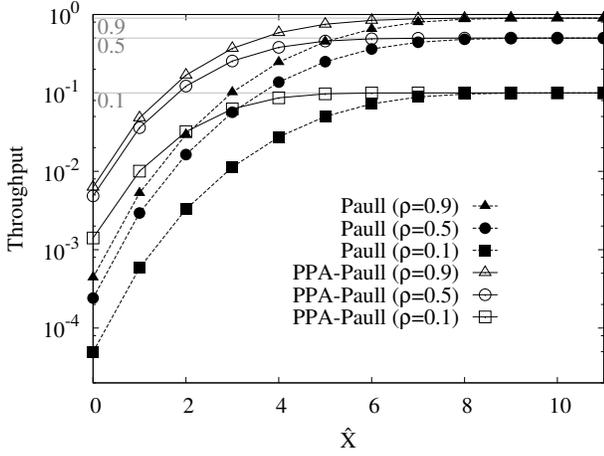


Fig. 10. Throughput under uniform traffic for $N = 64$.

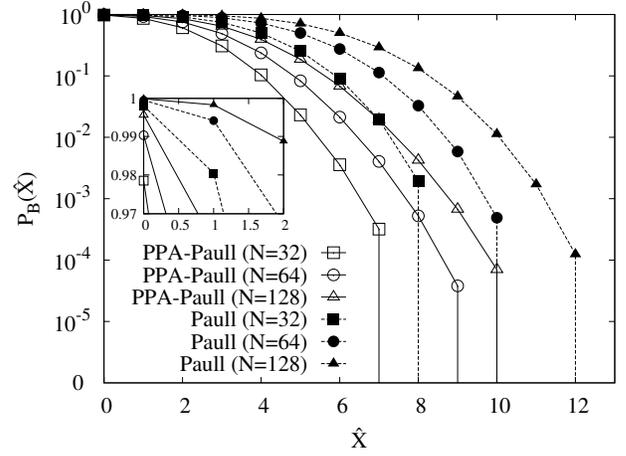


Fig. 12. Blocking probability in Benes networks with $\rho = 0.5$

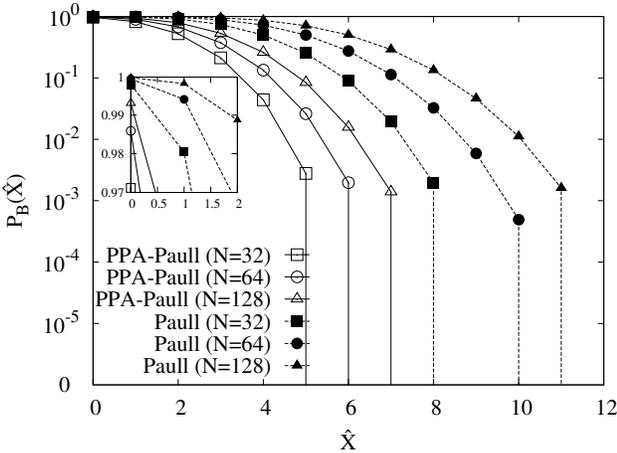


Fig. 11. Blocking probability in Benes networks with $\rho = 0.1$

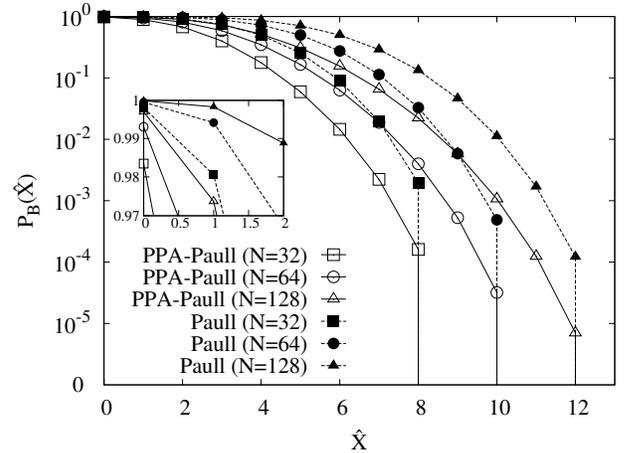


Fig. 13. Blocking probability in Benes networks with $\rho = 0.9$

is never blocked, for any permutation. In the figures, the throughput and the blocking probability are averaged across many timeslots.

Fig. 10 shows the throughput in function of the maximum degradation index \hat{X} for a 64×64 Benes network and different input loads: $\rho \in \{0.1, 0.5, 0.9\}$, corresponding to a lightly, medium and highly loaded network, respectively. For enough large \hat{X} , the throughput reaches its maximum value (0.1, 0.5 or 0.9 for each couple of curves), since the routing is not affected by the degradation index; in such case, all the algorithms behave the same. Note that the number of stages for the considered network is 11, hence larger values of \hat{X} are not affecting the routing. On the other side, smaller values of \hat{X} reduce the possibility of finding feasible paths; in the extreme case, the throughput approaches zero. In general, PPA-Paull achieves always a better throughput than Paull. When the input load ρ increases, the minimum \hat{X} to achieve the maximum throughput increases, since the routing is more constrained by a larger number of preliminary paths added during the current timeslot.

Fig. 11, Fig. 12 and Fig. 13 show the blocking probability in function of the maximum power penalty, each figure referring to a different value of input load. The smaller plot inside each figure details the blocking probability for low values of \hat{X} .

The total number of stages $S(N)$ in function of N are $S(32) = 9$, $S(64) = 11$ and $S(128) = 13$; whenever $\hat{X} > S(N)$, the blocking probability is zero by construction and the maximum throughput is achieved. On the contrary, when \hat{X} approaches zero, the routing is severely constrained by the degradation index: the blocking probability increases and the throughput tends to zero. Furthermore, as N increases, in all the figures the blocking probability increases due to the larger network depth. In general, the reduction in the blocking probability due to PPA-Paull with respect to Paull is very large, reaching more than two orders of magnitude in some cases.

To better understand such results, consider the case in which just one path must be connected (this event may happen for low input load). If we set $\hat{X} = 0$, there will be only one specific destination (among N possible ones) reachable by each input; the corresponding path will be found by PPA-Paull. Thus, at low load and under uniform traffic, we can expect $P_B^{\text{PPA-Paull}}(0) = 1 - 1/N$ (consistently with the values in the figure).

Now observe that, in a Benes network, there exist always $2^{\log_2 N - 1} = N/2$ different paths connecting any input to any output, since in the first $\log_2 N - 1$ stages there are always two output ports in each module that can be used to reach any destination, whereas in the last $\log_2 N$ stages

there exists just one output port in each module to reach the desired output. Given an input-output pair with a possible path compatible with $\hat{X} = 0$ (this pair is chosen with probability $1/N$ as shown above), the Paull algorithm will choose one random path among the $N/2$ available paths, but only one of them will be able to satisfy the constraint $\hat{X} = 0$. Hence, we can expect that $P_B^{\text{Paull}}(0) = 1 - 2/N^2$ (coherently with the values in the figure), which is larger than $P_B^{\text{PPA-Paull}}(0)$.

We now evaluate the maximum degradation index experienced by a single path computed by PPA-Paull. At each factorization level, PPA-Paull chooses the configuration of the first and the third stage to minimize the degradation index; in the case of a single path, the degradation index can increase by one at each factorization level. As a consequence, the maximum degradation index will be $\log_2 N - 1$, equal to the number of factorization levels. Hence, for low load, we expect that the blocking probability tends to zero when $\hat{X} \geq \log_2 N$. Indeed, Fig. 11 shows that the observed blocking probability goes to zero for $\hat{X} = 6, 7, 8$ when $N = 32, 64, 128$, which is very close to the bound found before.

VIII. CONCLUSIONS

In this paper we analyzed the scalability in terms of area and power penalty of different interconnection networks based on microring resonators. We described the basic 1×2 and 2×2 switching elements, we highlighted the asymmetric power loss of the different switching states, and we defined the degradation index. Then, we obtained the degradation index of crossbar, Benes and Clos networks. To achieve a better compromise between area (in terms of number of microrings) and degradation index, we propose (i) two architectures based on different combinations of the Benes and crossbar networks and (ii) the mirroring technique. Finally, we presented a simple variation of the classical Paull algorithm to setup input-output connections, and we investigated the corresponding improvement in terms of the degradation index. Given the promising results obtained in our studies, we believe that the role of microring resonators will become more and more relevant in future high capacity photonic interconnection networks.

REFERENCES

- [1] S. I. A. ITRS, "International technology roadmap for semiconductors," 2009, <http://www.itrs.net/Links/2009ITRS/Home2009.htm>.
- [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. Lee, A. Biberman, J. Chan, and K. Bergman, "High-performance modulators and switches for silicon photonic networks-on-chip," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 16, no. 1, pp. 6–22, 2010.
- [5] L. Tobing and P. Dumon, *Photonic Microresonator Research and Applications*, I. Chremmos, O. Schwelb, and N. Uzunoglu, Eds. Springer, 2010.
- [6] F. Xia, L. Sekaric, and Y. Vlasov, "Ultracompact optical buffers on a silicon chip," *Nature Photonics*, vol. 1, no. 1, pp. 65–71, 2006.
- [7] B. Lee, A. Biberman, N. Sherwood-Droz, C. Poitras, M. Lipson, and K. Bergman, "High-speed 2×2 switch for multiwavelength silicon-photonic networks-on-chip," *Journal of Lightwave Technology*, vol. 27, no. 14, pp. 2900–2907, 2009.

- [8] A. Bianco, D. Cuda, R. Gaudino, F. Neri, G. Gavilanes, and M. Petracca, "Scalability of optical interconnects based on microring resonators," *IEEE Photonic Technology Letters*, vol. 22, no. 15, pp. 1081–1083, July 2010.
- [9] R. Ji, L. Yang, L. Zhang, Y. Tian, J. Ding, H. Chen, Y. Lu, P. Zhou, and W. Zhu, "Microring-resonator-based four-port optical router for photonic networks-on-chip," *Optics Express*, vol. 19, no. 20, pp. 18 945–18 955, Sep 2011.
- [10] A. Bianco, D. Cuda, M. Garrich, R. Gaudino, G. Gavilanes, P. Giaccone, and F. Neri, "Optical interconnection architectures based on microring resonators," in *IEEE International Conference on Photonics in Switching (PS'09)*, Sept. 2009, pp. 1–2.
- [11] —, "Optical interconnection networks based on microring resonators," in *IEEE International Conference on Communications (ICC)*, May 2010, pp. 1–5.
- [12] J. Y. Hui, *Switching and traffic theory for integrated broadband networks*. Kluwer, 1990.
- [13] A. Biberman, N. Sherwood-Droz, B. Lee, M. Lipson, and K. Bergman, "Thermally active 4×4 non-blocking switch for networks-on-chip," in *IEEE 21st Annual Meeting of Lasers and Electro-Optics Society (LEOS)*, Nov 2008, pp. 370–371.
- [14] C. Li and A. Poon, "Silicon electro-optic switching based on coupled-microring resonators," in *Conference on Lasers and Electro-Optics (CLEO)*, May 2007.
- [15] A. Biberman, H. Lira, K. Padmaraju, N. Ophir, J. Chan, M. Lipson, and K. Bergman, "Broadband silicon photonic electro-optic switch for photonic interconnection networks," *IEEE Photonics Technology Letters*, no. 99, pp. 1–1, 2011.
- [16] 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.
- [17] P. Sanchis, P. Villalba, F. Cuesta, A. Håkansson, A. Griol, J. Galán, A. Brimont, and J. Marti, "Highly efficient crossing structure for silicon-on-insulator waveguides," *Optics letters*, vol. 34, no. 18, pp. 2760–2762, 2009.
- [18] W. Bogaerts, P. Dumon, D. Thourhout, and R. Baets, "Low-loss, low-cross-talk crossings for silicon-on-insulator nanophotonic waveguides," *Optics letters*, vol. 32, no. 19, pp. 2801–2803, 2007.
- [19] J. Van Campenhout, W. Green, S. Assefa, and Y. Vlasov, "Low-power, 2×2 silicon electro-optic switch with 110-nm bandwidth for broadband reconfigurable optical networks," *Optics Express*, vol. 17, no. 26, pp. 24 020–24 029, 2009.
- [20] P. Dong, S. Liao, H. Liang, R. Shaffiha, D. Feng, G. Li, X. Zheng, A. Krishnamoorthy, and M. Asghari, "Submilliwatt, ultrafast and broadband electro-optic silicon switches," *Optics Express*, vol. 18, no. 24, pp. 25 225–25 231, 2010.
- [21] Y. Kokubun, Y. Hatakeyama, M. Ogata, S. Suzuki, and N. Zazizen, "Fabrication technologies for vertically coupled microring resonator with multilevel crossing busline and ultracompacting radius," *IEEE Journal of Selected Topics in Quantum Electronics*, vol. 11, no. 1, pp. 4–10, Feb. 2005.

PLACE
PHOTO
HERE

Andrea Bianco (M'98-SM'09) received the Ph.D. degree in telecommunications engineering from Politecnico di Torino, Italy, in 1993. He is a Full Professor and Vice-Head of the Dipartimento di Elettronica e Telecomunicazioni of Politecnico di Torino. He has coauthored over 150 papers published in international journals and presented in leading international conferences in the area of telecommunication networks. His current main research interests are in the fields of protocols

and architectures for all-optical networks and switch architectures for high-speed networks. Dr. Bianco is Associate Editor for the Elsevier Computer Communications Journal and has been Guest or Co-Guest Editor of several special issues in international journals, including the IEEE Communications Magazine and Computer Networks.

PLACE
PHOTO
HERE

Davide Cuda (M'01) received his Dr.Ing. and his Ph.D. degrees in Information Technologies at Politecnico di Torino, Torino, Italy, in 2005 and 2009, respectively. He is currently a PostDoc researcher at France Telecom R&D. From August 2007 to August 2008 he visited the Networks Research Laboratory led by Prof. Biswanath Mukherjee at University of California, Davis. His main area of interests are the design of scheduling policies for high-performance routers and all-optical networks

and all-optical switching architectures.

PLACE
PHOTO
HERE

Guido Gavilanes Castillo graduated from the Universidad del Cauca (Colombia) in Electronics and Telecommunication Engineering in May 2002. During the last year of his studies he was member of Telematics Engineering R+D Group of Universidad del Cauca. He received his Master Degree in 2006 on Optical Communication and Photonic Technologies from Politecnico di Torino. In 2010 he obtained his PhD degree working with the Telecommunication Networks Group, Politecnico di Torino. He studied architectures and algorithms applied to optical switching systems mainly using passive optical components, such as AWGs, and Optical Microring Resonators. Since May 2011 he has been working at the Networking Laboratory at the Istituto Superiore Mario Boella, Turin, Italy, investigating optical network architectures for future Networks on Chip.

He studied architectures and algorithms applied to optical switching systems mainly using passive optical components, such as AWGs, and Optical Microring Resonators. Since May 2011 he has been working at the Networking Laboratory at the Istituto Superiore Mario Boella, Turin, Italy, investigating optical network architectures for future Networks on Chip.

PLACE
PHOTO
HERE

Miquel Garrich (S'11) was born in Barcelona, Spain. He received the Dr.Ing. in telecommunications engineering from Universitat Politecnica de Catalunya in Barcelona, Spain, in October 2009. From September 2008 to October 2009 he was student visitor for his master thesis at the Politecnico di Torino, Italy. Currently, he is working toward the Ph.D. degree in the Dipartimento di Elettronica e Telecomunicazioni of Politecnico di Torino,

Italy. From July 2011 to April 2012 he was Academic Visitor at the School of Computer Science and Electronic Engineering, University of Essex, UK. His current research interests include switching interconnection architectures, optical technology and high performance optical networks.

PLACE
PHOTO
HERE

Paolo Giaccone (M'98) received the Dr.Ing. and Ph.D. degrees in telecommunications engineering from Politecnico di Torino, Italy, in 1998 and 2001, respectively. He is currently an Assistant Professor in the Dipartimento di Elettronica e Telecomunicazioni, Politecnico di Torino. During the summer of 1998, he was with the High Speed Networks Research Group, Lucent Technology-Bell Labs, Holmdel, NJ. During 2000-2001 and in 2002 he was with the Information Systems Net-

working Lab, Electrical Engineering Dept., Stanford University, Stanford, CA. His main area of interest is the design of network algorithms, the theory of interconnection networks, and the performance evaluation of telecommunication networks through simulative and analytical methods.