



Energy-Efficiency Comparison of Multi-Layer Deposited Nanophotonic Crossbar Interconnects

Hui Li, Sébastien Le Beux, Martha Johanna Sepulveda Florez, Ian O'Connor

► To cite this version:

Hui Li, Sébastien Le Beux, Martha Johanna Sepulveda Florez, Ian O'Connor. Energy-Efficiency Comparison of Multi-Layer Deposited Nanophotonic Crossbar Interconnects. ACM Journal on Emerging Technologies in Computing Systems, 2017, XX. hal-01508192

HAL Id: hal-01508192

<https://inria.hal.science/hal-01508192>

Submitted on 14 Apr 2017

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Energy-Efficiency Comparison of Multi-Layer Deposited Nanophotonic Crossbar Interconnects¹

Hui Li, Lyon Institute of Nanotechnology, Ecole Centrale de Lyon
Sébastien Le Beux, Lyon Institute of Nanotechnology, Ecole Centrale de Lyon
Martha Johanna Sepulveda, Laboratory CAIRN/IRISA INRIA Rennes
Ian O'Connor, Lyon Institute of Nanotechnology, Ecole Centrale de Lyon

Single-layer optical crossbar interconnections based on Wavelength Division Multiplexing (WDM) stand among other nanophotonic interconnects by their low latency and low power. However, such architectures suffer from a poor scalability due to losses induced by long propagation distances on waveguides and waveguide crossings. Multi-layer deposited silicon technology allows the stacking of optical layers which are connected by means of Optical Vertical Coupler. This allows significant reduction in the optical losses, which contributes to improve the interconnect scalability, but also leads to new challenges related to network designs and layouts. In this paper, we investigate the design of optical crossbars using multi-layer silicon deposited technology. We propose implementations for Ring, Matrix, λ -router and Snake based topologies. Layouts avoiding waveguide crossings are compared to those minimizing the waveguide length according to worst-case and average losses. The laser output power is estimated from the losses, which allows evaluating the energy efficiency improvement induced by multi-layer technology over traditional planar implementations (33% on average). Finally, networks comparison has been carried out and the results show that the ring topology leads to a 43% reduction in the laser output power.

CCS Concepts: • Networks-Network architectures • Networks-Network performance evaluation • Networks-Network on chip

Additional Key Words and Phrases: Optical Network on Chip, crossbar, multi-layer, optical loss, energy efficient

ACM Reference Format:

LI, H., LE BEUX, S., JOHANNA SEPULVEDA, M., and O'CONNOR, I., 2017. Energy-Efficiency Comparison of Multi-Layer Deposited Nanophotonic Crossbar Interconnects. *ACM J. Emerg. Technol. Comput. Syst.*, xx, Article xx (XX 2017), xx pages. DOI: xxxxxxxxxxxxxxxx

1. INTRODUCTION

Inter-core communication is currently a major bottleneck to achieve high performance Multi-Processors System-on-Chip (MPSoCs). 3D die stacking technology appeared as a promising

¹ A preliminary version of this article was presented at the 20th Asia and South Pacific Design Automation Conference (ASP-DAC) 2015.

Author's addresses: H. LI, Lyon Institute of Nanotechnology, Ecole Centrale de Lyon, France; S. LE BEUX (corresponding author), Lyon Institute of Nanotechnology, Ecole Centrale de Lyon, France, email: sebastien.le-beux@ec-lyon.fr; M. JOHANNA SEPULVEDA, Laboratory CAIRN/IRISA INRIA Rennes, France; I. O'CONNOR, Lyon Institute of Nanotechnology, Ecole Centrale de Lyon, France.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. To copy otherwise, distribute, republish, or post, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

© 2017 Copyright held by the owner/author(s). Publication rights licensed to ACM. 123-4567-24-567/08/ART6.\$15.00
DOI: xxxxxxxxxxxxxxxx

solution to overcome this bottleneck by reducing the distance between the cores. In such architecture, intra-layer communication is usually carried out by planar Electrical Networks-on-Chip (ENoCs) while inter-layer communications rely on Through-Silicon-Vias (TSVs). The integration of heterogeneous technologies now allows new interconnect options to be explored, such as silicon photonics [1], which has the potential to improve communication latency and bandwidth [2][3][4]. Silicon photonic interconnect is traditionally implemented on an optical layer integrating laser sources, Microring Resonators (MRs), photodetectors and waveguides.

Among the proposed optical interconnects, wavelength-routing based solutions stand by their low latency and low power, since they use passive MRs and do not require any arbitration [4][5]. In such networks, the communication between a source IP core and a destination IP core is carried out through one wavelength or a set of wavelengths by using Wavelength Division Multiplexing (WDM). Despite these advantages, existing optical crossbars show different tradeoffs between design complexity and energy efficiency. Moreover, their structures are highly penalized by a lack of scalability due to propagation losses, waveguide crossing losses, switching losses and drop losses. Waveguide crossing is a major source of losses and it can reach 0.2dB per crossing [6].

Emerging design technology based on multi-layer deposited silicon enables the efficient stacking of optical layers [7][8]. They rely on Optical Vertical Couplers (OVCs) implemented by using inverse tapers [9], multimode interference (MMI) [10], Microring Resonator (MR) [11][12] or grating-assisted vertical coupler [13][14]. Multi-layer deposited silicon contributes to reduce the number of waveguide crossings, but leads to new losses related to inter-layer coupling. Design trade-off thus needs to be explored, in order to improve the optical interconnect energy efficiency.

In this paper, we investigate the design of optical crossbars using multi-layer silicon deposited technology. We propose implementations for ring [5], Matrix [15], λ -router [4] and Snake [16] based topologies. Layouts avoiding waveguide crossings are compared to those minimizing the waveguide length according to the worst-case and average losses. The laser output power is estimated from the losses, which allows evaluating the energy efficiency improvement induced by multi-layer technology over traditional planar implementations.

The paper is structured as follows. Section 2 presents the related work. Section 3 presents the considered 3D architecture model. Section 4 and Section 5 present the proposed multi-layer implementations of Ring, Matrix, λ -router and Snake crossbars. Section 6 gives the evaluation and exploration results. Finally, Section 7 concludes the paper.

2. RELATED WORK

The design of nanophotonic interconnects has been thoroughly investigated in the literature. Among the numerous proposed solutions, Flexishare [24] is one of the most flexible interconnects since it relies on a reservation assisted MWMR (Multiple Writer Multiple Readers) communication scheme. However, opening a communication channel in such a network requires arbitration on both writers and receivers sides. This leads to latency and create contention on the interfaces. Central controllers have been proposed to accelerate the communication channel management [25] but such approaches are not scalable to large-scale systems. Wavelength routed ONoCs (WRONoCs) crossbars do not suffer from such latency and contention since no arbitration is needed: point-to-point communication channels between all the interfaces are permanently opened. Matrix [15], λ -router [4], Snake [16] and ORNoC [5] are WRONoCs that rely on passive MRs. Each network exhibits different characteristics such as number of optical resources, insertion loss and scalability. Although the design of optical crossbar interconnection has become popular, only the works [16] and [20] compare them under a given connectivity scenario. In [16], the authors compare λ -router, Snake and ORNoC, which

perform the data exchange among the processing and storage components of a many-core system. As a result, they show that ORNoC achieves higher energy efficiency. In our prior work [20], we compare the above-mentioned optical crossbars for different system sizes according to the worst-case loss metric. All these previous works only address single-layer based WRONoC implementations. In our previous work [26], a very first study has been carried out to compare WRONoCs implemented using this technology. This paper further investigates their implementations by exploring technology-related design parameters. We also propose an algorithm optimizing the design of ring topology based WRONoC.

3D-IC allows the stacking of heterogeneous layers that are connected using Through Silicon Vias (TSVs) [17]. In [18], an optical layer implementing lasers is stacked on top of an optical layer on which a network is implemented. Recently, multi-layer deposited silicon has been introduced as a key technology allowing the stacking of passive optical devices [10]. This helps reducing waveguide crossing losses (typical value ranges from 0.05dB [27] to 0.2dB [6]), which makes this technology highly suitable to optimize network suffering from a higher number of crossings [19]. In this context, a four-layer static optical crossbar has been proposed in [21] and a reconfigurable version of this network has been proposed in [22][23]. These approaches are complementary to our work since the networks could be further improved by using multi-layer deposited silicon.

3. ARCHITECTURE MODEL

In this section, we first present the architecture of a multi-layer optical interconnection. The second subsection describes the multi-layer technology. The third subsection describes the optical worst-case and average losses models used to evaluate the minimum laser output power.

3.1 Multi-layer architecture overview

The considered 3D architecture is composed of an electrical layer and two optical layers. Figure 1 illustrates an architecture example for a 4x4 cores and by assuming a ring topology for the optical interconnect. The electrical layer is composed of IP cores which process and store data. The IP cores are arranged into an $N \times N$ mesh, with N an even number. The data among the cores are exchanged through a WRONoC implemented using the optical layers. Electrical layer is connected to the optical layer by means of TSV [28], a set of conductive nails that extend out the back-side of a thinned-down die. The optical layers are composed of on-chip lasers such as Vertical-Cavity Surface Emitting Lasers (VCSELs) [29], Microring Resonators (MRs), photodetectors and waveguides. We assume the use of on-chip lasers since they don't lead to the use of power waveguides as with off-chip lasers [30], which contribute to reduce the number of waveguide crossings, and thus the total losses in the communication paths.

In WRONoC, the communications depend on the signals wavelength: each wavelength is assigned to a source/destination cores pair at design time. This leads to low latency communication since no arbitration is needed. The optical devices are gathered into the transmitter and the receiver parts of optical network interfaces (ONIs). The transmitter allows optical signals to be emitted with direct modulation at different wavelengths and coupled into waveguides. In the receiver part, the optical signals are ejected from the waveguides and are redirected to photodetector for O/E conversion. The routing of the signals is achieved according to the signal wavelength λ_s and the resonating wavelengths of crossed MRs. In this work, we consider a fully connected optical crossbar. It interconnects cores by means of $(N^2-1) \times N^2$ laser sources, $(N^2-1) \times N^2$ photodetectors, and $(N^2-1) \times N^2$ passive MRs. N^2 is the total number of IP cores by considering $N \times N$ mesh.




Figure 1: The optical crossbar (ring topology in the example) is implemented in the optical layers and it interconnects IP cores in the electrical layer.

The minimum laser output power depends on the total loss experienced by the optical signals from the source ONI to the destination ONI. The higher the losses, the higher the required laser output power, i.e., the lower the energy efficiency. It is worth noticing that the received optical power should be high enough to reach the SNR requirements for a given target BER, which is out of the scope of the paper but has been investigated in [31]. Reducing the losses is thus mandatory to improve the overall system energy efficiency. Among the sources of losses, the most significant ones are those related to the signal propagation in the waveguides, the waveguide crossing and the MR drop. In this paper, we investigate losses reduction achievable using multi-layer silicon deposited technology.

3.2 Multi-layer deposited silicon technology

Multi-layer deposited silicon allows optical interconnects to be improved by stacking optical layers [27][32]. Indeed, waveguide crossings can be avoided, as illustrated in Figure 2-a. In the figure, red and blue colors represent waveguides implemented in the first and second layer, respectively. Figure 3-a illustrates the top-view of a waveguide crossing implemented with a single-layer (which leads to 0.05-0.2dB loss) and with two layers (nearly 0dB losses when waveguides are placed orthogonally with an appropriate vertical gap [33]).




Figure 2: a) 3D view of waveguide crossing in different layers, and Optical Vertical Coupler (OVC) based on: b) inverse tapers and c) MMI.

Optical signals propagate from one layer to another by using OVCs which can be designed based on inverse tapers (Figure 2-b) or MMI (Figure 2-c). The coupling efficiency, given by the power ratio between OP_{out} and OP_{in} , depends on the physical dimensions of the waveguide (i.e., height H and width W), the properties of the taper (e.g., type of material) and their location on the circuit (i.e., vertical gap g , tips longitudinal overlapping L and taper angle θ) [9]. Figure 3-b illustrates the top-view of a waveguide designed with single-layer and multi-layer technologies.




Figure 3: Implementations with one layer and two layers of: a) waveguide crossing, b) waveguide, c) MR, and d) PSE. The insertion loss values given in the figure are extracted from [27].

The propagation losses can also be reduced by considering, for instance, silicon nitride (Si_3N_4) (layer 2) deposited on top of a standard silicon on insulator (SOI) (layer 1) [34][35]. For layer 1, we assume a crystalline silicon (c-Si) waveguide, with a cross-section dimension of $500\text{nm} \times 220\text{nm}$ ($W \times H$) and a refractive index (n_{Si}) of 3.45. A reported propagation loss is 2.85dB/cm [34] and it has been reduced to 0.5dB/cm [27]. For layer 2, we assume a CMOS-compatible silicon nitride (Si_3N_4) waveguide, with a cross-section dimension of $1000\text{nm} \times 400\text{nm}$ ($W \times H$) and a refractive index ($n_{\text{Si}_3\text{N}_4}$) of 2. Reported propagation loss is 1.3dB/cm [34] around 1550nm and optimized implementations allow reducing the loss to 0.1dB/cm [27]. By using silicon dioxide (SiO_2) as the cladding ($n_{\text{SiO}_2}=1.5$), high confinement of the optical signal and sharp bending radius are achieved. In other words, once an optical signal reaches layer 2, it will experience lower propagation losses compared to signal propagating on layer 1. However, reaching layer 2 is possible only by crossing OVCs, which leads to additional losses L_{OVC} (e.g., 0.2dB and 0.1dB reported in [34] and [27], respectively). 3D implementation of photonic devices is also possible. The structure of the networks being independent from the implementation technology, the number of devices (lasers, photodetectors and MRs) is the same for single-layer and two-layer implementations. The only potential footprint overhead comes from the OVCs,

which is designed by overlapping waveguides located on both layers. Figure 2-b and -c represent the overlap of a distance L . However, L can be as small as 20um [9], which we neglect considering that the whole die size (typically 2cm×2cm) is dedicated to the implementation of the crossbar. The main overhead is related to the additional fabrication complexity and higher design costs induced by stacking multiple silicon layers, which is not evaluated in the paper.

The switching operation in MRs (Figure 3-c) depends on the signal wavelength (λ_s) and MR resonant wavelength (λ_x). When λ_s is equal to λ_x , the signal coming from the horizontal waveguide will couple into the MR, thus being redirected to the vertical waveguide. When λ_s is different from λ_x , the signal continues propagating in the horizontal waveguide. A Photonic Switching Elements (PSEs) [3][4] is composed of 2 crossing waveguides and 2 MRs with a same resonant wavelength λ_x (Figure 3-d). Depending on their wavelength λ_s , an optical signal is routed as follow: for $\lambda_s = \lambda_x$, a resonance occurs in the MRs and the signal is redirected on the other waveguide; otherwise, no resonance occurs and the signal continues propagating on the same waveguide. PSEs are the basic blocks of multistage networks (e.g. Snake and λ -router), which require 2 inputs and 2 inputs in switching structures. The routing of the signals in the network depends on i) the resonant wavelengths of the PSEs and ii) the way how PSEs are connected to each other. MRs [27] and PSEs can be also efficiently implemented by means of the multi-layer technology since waveguide crossing is avoided, as illustrated in Figure 3-c and Figure 3-d.

3.3 Models for worst-case and average losses

The worst-case and average losses are key metrics to measure the energy efficiency in optical interconnects since they allow estimating the minimum required laser output power. Figure 3 presents the losses parameters we assume and the total loss along an optical path L_{total} (after the coupling of the optical signal emitted by the laser into the waveguide) is given in equation (1), which is an extension of the model proposed in [20]. L_{total} depends on: i) the total propagation loss in the waveguide $L_{propagation}$, given by equation (1-a) and represented in Figure 3-b; ii) the total loss due to the effective number of waveguide crossings $L_{crossing}$ given by equation (1-b) and represented in Figure 3-a; iii) the total drop loss L_{drop} given by equation (1-c) and represented in Figure 3-c and -d; iv) the coupler loss L_{OVC} given by equation (1-d) and represented in Figure 3-b; v) the total through loss $L_{through}$ given by equation (1-e) when a signal passes by a non-resonant MR; and vi) the waveguide bending loss $L_{bending}$. In this work, we assume $L_{bending}$ (e.g., $P_{bending} = 0.005\text{dB}/90^\circ$ [27]) and $L_{through}$ is neglected. We also assume negligible crosstalk between waveguides, which can be obtained by considering a 5μm distance between parallel waveguides. Indeed, for 5mm parallel waveguides assuming 500nm×220nm ($W \times H$), the power coupling between the waveguides will be lower than -40dB, when the gap side by side is 3 μm or more [36]. The loss induced by the fabrication process variation is not considered in this work. The parameters used in the formulation are detailed in Table 1 and Table 2.

$$L_{total}^{dB} = L_{propagation}^{dB} + L_{crossing}^{dB} + L_{drop}^{dB} + L_{through}^{dB} + L_{bending}^{dB} + L_{OVC}^{dB} \quad (1)$$

$$L_{propagation} = P_{propagation,1} \times l_{s-d,1} + P_{propagation,2} \times l_{s-d,2} \quad (1-a)$$

$$L_{crossing} = P_{crossing} \times N_{crossing} \quad (1-b)$$

$$L_{drop} = P_{drop,1} \times N_{drop,1} + P_{drop,2} \times N_{drop,2} \quad (1-c)$$

$$L_{OVC} = P_{OVC} \times N_{OVC} \quad (1-d)$$

$$L_{through} = P_{through} \times N_{through} \quad (1-e)$$

Table 1 Insertion Loss Parameters

Parameter	Description
$P_{propagation,1} (dB/cm)$	Intrinsic propagation loss of waveguide in layer 1
$P_{propagation,2} (dB/cm)$	Intrinsic propagation loss of waveguide in layer 2
$P_{crossing} (dB)$	Waveguide crossing loss
$P_{drop,1} (dB)$	Drop loss in the same layer
$P_{drop,2} (dB)$	Drop loss in MR and PSE in different layers
$P_{OVC} (dB)$	Vertical coupling loss (in OVC)
$P_{through} (dB)$	Through loss when a signal crosses a non-resonant MR

Table 2 Network Implementation Characteristics

Parameter	Description
$l_{s-d,1}$	Waveguide length between a source and a destination in layer 1
$l_{s-d,2}$	Waveguide length between a source and a destination in layer 2
$N_{crossing}$	Number of waveguide crossings
$N_{drop,1}$	Number of intra-layer drop operations
$N_{drop,2}$	Number of inter-layer drop operations
N_{OVC}	Number of vertical couplers along a path
$N_{through}$	Number of MRs crossed by signals

From the total loss along a communication path between any pair of source/destination IP cores (shown in equation (1)), the worst-case loss (L_{wc}) and the average loss (L_{avg}) are estimated by using equation (2) and (3). In these equations, L is the set of total losses (i.e., L_{total}) for all the communication paths in the network. This model is generic and can be used for both single-layer and multi-layer implementations.

$$L_{wc} = \text{Maximum}(L) \quad (2)$$

$$L_{avg} = \text{Average}(L) \quad (3)$$

From the worst-case loss (L_{wc}) and the receiver sensitivity ($OP_{sensitivity}$), the minimum laser output power (OP_{min_laser}) required for a given BER can be obtained as following:

$$OP_{min_laser}^{dBm} = L_{wc}^{dB} + OP_{sensitivity}^{dBm}$$

In the results section, WRONoCs will be compared based on their required minimum laser output power.

4. MULTI-LAYER OPTICAL RING CROSSBAR

In this section, we present ORNoC_{ML}, a ring-based optical crossbar implemented with multi-layer deposited silicon technology. We first present the topology and connectivity of the optical crossbar. Then, the design method for ORNoC_{ML} is presented.

4.1 Multi-Layer Implementation of Ring Based WRONoC

ORNoC is a ring based optical crossbar [5][37] illustrated in the left-hand side of Figure 4. The main feature of ORNoC is the absence of waveguide crossings, which is possible due to the serpentine layout and the use of on-chip lasers. In the figure, solid and dot lines represent the clockwise (C) and counter-clockwise (CC) directions for signal propagations, respectively. In ORNoC, the waveguides are segmented and ONIs define the limits of the segments (in our architecture, we assume that each ONI is linked to a given IP). ONIs are designed to i) inject optical signals into a segment, ii) eject received optical signals from segment and iii) let optical signals propagating through the waveguide (i.e. propagating from a segment to another). Signal injection and ejection are achieved using MRs, as detailed in [5][37]. Hence, the number of segments crossed by an optical signal depends on its wavelength λ_s and the resonant wavelengths of the MRs located on the waveguide. Once the signal is ejected from the waveguide, the corresponding wavelength is free and can be used for another communication. For instance, assuming λ_0 is used for $IP_1 \rightarrow IP_2$ communication, it can also be used for $IP_2 \rightarrow IP_3$. Obviously, a same wavelength cannot be used on a same segment for two different communications. Hence, $IP_1 \rightarrow IP_3$ communication will require another wavelength, for instance λ_1 . Furthermore, multiple waveguides can be used to transmit optical signal in C and CC directions.




Figure 4: Optical crossbars ORNoC and ORNoC_{ML}: a) topology to interconnect 9 IP cores and b) layout for 4x4 IP cores.

ORNoC_{ML} is the multi-layer implementation of ORNoC and is illustrated in the right-hand side of Figure 4. It implements a second set of rings located on the second layer, with the aim to improve the connectivity between the IP cores thanks to reduced losses. Red and blue colors are used to represent waveguides located in the first and second layer, respectively. The ring layouts

in the second layer are rotated by 90° compared to the first layer layout. Since the additional waveguides are located in a different layer, the propagation of signal does not suffer from any additional waveguide crossing loss.

The following illustrates the advantages of ORNoC_{ML} over ORNoC, assuming the same propagation loss value for both layers for the sake of clarity. The left-hand side of Figure 4-b shows the single-layer ORNoC layout for 4×4 cores. In order to perform the communication with the lowest $L_{\text{propagation}}$ between the $IP_1 \rightarrow IP_9$ and $IP_4 \rightarrow IP_2$, the C and CC directions are employed, respectively. Note that the response communications (i.e., $IP_9 \rightarrow IP_1$ and $IP_2 \rightarrow IP_4$) will be performed in opposite directions, i.e., by using CC and C. $IP_1 \rightarrow IP_9$ is one of the communication paths that experience the most losses. Single-layer ORNoC implementation requires the crossing of 7 intermediate interfaces. By considering a mesh distribution of the interfaces and a distance d between neighboring IP cores, the total propagation distance is thus $8d$. In order to reduce this distance, dedicated waveguide for $IP_1 \rightarrow IP_9$ can be integrated in the same layer (e.g., $IP_1 \rightarrow IP_2 \rightarrow IP_3 \rightarrow IP_6 \rightarrow IP_9$). However, this will: i) introduce waveguide crossings; and ii) affect the regularity, thus leading into a less scalable network. With ORNoC_{ML} (right-hand side of Figure 4-b), $IP_1 \rightarrow IP_9$ and $IP_9 \rightarrow IP_1$ are implemented on the second layer by using C and CC directions since the propagation distance is shorter than that in the first layer. Communications between $IP_2 \rightarrow IP_4$ and $IP_4 \rightarrow IP_2$ are still implemented in the first layer. Hence, ORNoC_{ML} avoids the introduction of additional waveguide crossings and reduces the propagation distance, while keeping the layout regular.

Figure 5 represents the worst case and the average number of crossed interfaces (which lead to the worst case and the average losses respectively) for ORNoC and ORNoC_{ML}. As illustrated in Figure 5-a, the second layer doesn't allow improving the worst-case distance, which is due to the serpentine layout of both networks. However, significant reduction in the average number of crossed interfaces is achieved, which allows global improvement of the network energy efficiency. Furthermore, since the reduction in the average distance increases with the network size, the second layer contributes to the ONoC scalability.




Figure 5: Number of crossed intermediate interfaces for ORNoC and ORNoC_{ML} under 2×2 , 4×4 , 6×6 and 8×8 network sizes: a) worst-case and b) average case.

Regarding the second layer, many design options are actually possible and, as a first constraint, a similar layout is kept on the two layers in order to allow partial reuse of the backend optimization results (e.g. distance between waveguides and bending curves), which contributes to reduce the fabrication cost. We thus selected the serpentine layout for the second layer. Several design options are compared, including 90° and 270° rotations of the ring. 90° rotation demonstrated the lowest losses, thus we focus on this design option in the paper. Furthermore, these configurations are those leading to perpendicular crossing [32] of the

waveguides located on different layers: the overlapping distance is the smallest possible, which minimize the vertical coupling of the waveguide and hence leads to the lowest losses. Considering other angles (e.g. 45°) helps reducing the maximal distance between IPs located on a same diagonal, but this comes at the cost of a higher overlap between crossing waveguides located on different layers. In other words, the crossing structure will lead to optical power leakage from a waveguide to another. There is thus a trade-off between the propagation loss and the losses for non- 90° crossing structure, which is out of the scope of the paper.




Figure 6: An Optical Network Interface.

Figure 6 illustrates a layout example for an ONI in ORNoC_{ML}. The MRs and photodetectors are responsible for receiving optical signals and on-chip laser sources are used for emitting optical signals. The waveguides in red and blue allow the propagation of optical signals in the first and second layer, respectively. In this example, a single waveguide is considered. However, multiple waveguides can be regularly implemented without any waveguide crossing by applying the layout guidelines from [5]. Communications occurring on layer 1 will be achieved as in the single layer implementation of ORNoC. The signal propagating on the second layer will cross two OVCs: the first vertical coupling will occur right after their emission by the laser (i.e., layer 1 \rightarrow layer 2) and the second coupling will occur just before their reception by the photodetector (i.e., layer 2 \rightarrow layer 1).

All the on-chip laser sources and photodetectors are located in layer 1, which are turned on only when communications occur. In our work, we consider the PCM-VCSELs (illustrated in Figure 7-a) as on-chip laser sources. They rely on a double set of Si/SiO₂ photonic crystal mirrors (PCMs). PCM-VCSELs are considered due to their micrometer-scale layer thickness (thinner than VCSELs using DBR), their broadband reflectivity, full control over the cavity modal and polarization emission features [38]. Moreover, PCM-VCSELs are CMOS compatible. The fabrication employs standard CMOS pilot line processing tools and high-yield full-wafer bonding of group III-V alloys on silicon [38]. Coupling the vertical light from VCSEL into a horizontal waveguide can be achieved by using a taper located on the layer of the top PCM and the waveguide (as shown in Figure 7-b). We assume an 80% coupling efficiency, which is slightly pessimistic compared to the 85% simulated in [39].




Figure 7: PCM-VCSEL: a) 3D view extracted from [38] and b) cross-section view including the taper.

4.2 Design method

ORNoC_{ML} is designed following a two-step methodology. First, each communication is assigned to a ring (i.e., layer/direction couple) minimizing the total loss. Then, for each ring, wavelengths are assigned to the communication following an iterative algorithm. The following details the method.

4.2.1 First step: ring assignment

In the crossbars we assume, the entire possible source to destination communications schemes are carried out. In the first step, we allocate, for each communication, the optical path with the lowest propagation losses in the rings (i.e., layer 1/2 and direction C/CC). For this purpose, four distance matrixes are computed, each for one possible ring implementation (i.e., layer 1 clockwise, layer 1 counter-clockwise, layer 2 clockwise and layer 2 counter-clockwise). Each communication is assigned to the layer-direction couple showing the lowest loss.

Figure 8 illustrates an excerpt of the ring assignment for the 4×4 architecture illustrated in Figure 4, assuming the same propagation loss value for layer 1 and layer 2. Source and destination IP cores are represented in column and row, respectively. In this example, all the communications between IP₁, IP₂, IP₃ and IP₄ use the rings located on the first layer. Some communications will use C ring (e.g. IP₁→IP₂) and the others CC (e.g., IP₂→IP₁). As another example, IP₁→IP₉ and IP₉→IP₁ are implemented by using layer 2 in C and CC directions, respectively. In case the distances on layer 1 and layer 2 are the same, layer 1 is used in order to avoid vertical couplers. The wavelength assignment in each ring is achieved in the second step.

D \ S	IP ₁	IP ₂	IP ₃	IP ₄	...	IP ₉	IP ₁₀	IP ₁₁	...
IP ₁	-	C	C	C		C	CC	CC	
IP ₂	CC	-	C	C		C	C	CC	
IP ₃	CC	CC	-	C		CC	C	C	
IP ₄	CC	CC	CC	-		CC	C	C	
...									
IP ₉	CC	CC	C	C		-	C	C	
IP ₁₀	C	CC	CC	CC		CC	-	C	
IP ₁₁	C	C	CC	CC		CC	CC	-	
...									

Figure 8: Ring assignment matrix for 4×4 cores. Red and blue colors represent layer 1 and layer 2, respectively. C and CC denote clockwise and counter-clockwise directions, separately.

4.2.2 Second step: wavelength assignment algorithm

The design of ORNoC_{ML} requires careful wavelength assignment between cores in order to minimize the number of wavelengths and the number of waveguides. For this purpose, an

algorithm is executed for each of the 4 rings. Inputs of the algorithm are a list of communications to be assigned and the maximum number of wavelengths per waveguide. From ring assignment obtained in the first step, the wavelengths are assigned as follow. For each ring, initial IP (source IP_i), waveguide wg_i and wavelength λ_i are first defined. Then, λ_i is assigned to the shortest optical path from IP_i , which allows reaching an intermediate destination IP_d . The operation is repeated from IP_d , until IP_i is reached (i.e., wavelength λ_i has been assigned on all the segments of the waveguide). Another wavelength (λ_{i+1}) is used and the assignment process is repeated until the wavelength has been assigned to all the communications starting from the initial IP source. Then, a new wavelength is used and the process restarts from the following IP (IP_{i+1}), etc. If the number of wavelengths reaches the maximum allowed per waveguide, a waveguide is added (wg_{i+1}) and the algorithm continues its execution but from the initial wavelength λ_i . There is no limitation in the number of waveguides and, for symmetry purpose, bidirectional communications (i.e., communications occurring on a same layer but in opposite directions) are implemented with the same wavelength but in different waveguides.

Figure 9 illustrates the main steps of the algorithm for 5 IP cores, assuming a maximum of two wavelengths per waveguide. Starting from IP_1 , wavelength λ_1 is assigned in a first waveguide wg_1 to reach the closest intermediate destination ($IP_1 \rightarrow IP_2$ arrow in Figure 9-a). The process repeats with the same wavelength until the initial core is reached (in Figure 9-b, λ_1 is assigned to $IP_2 \rightarrow IP_3$, $IP_3 \rightarrow IP_4$, $IP_4 \rightarrow IP_5$, and $IP_5 \rightarrow IP_1$). Then, λ_2 is selected and the process starts again to reach the closest destination for which no wavelength has been assigned ($IP_1 \rightarrow IP_3$ in Figure 9-b). Once the number of wavelengths per waveguide is reached, a waveguide is added and the algorithm continues iterating with λ_1 ($IP_2 \rightarrow IP_4$ in Figure 9-c) until a waveguide and a wavelength has been assigned to all the communications in the matrix (Figure 9-d).

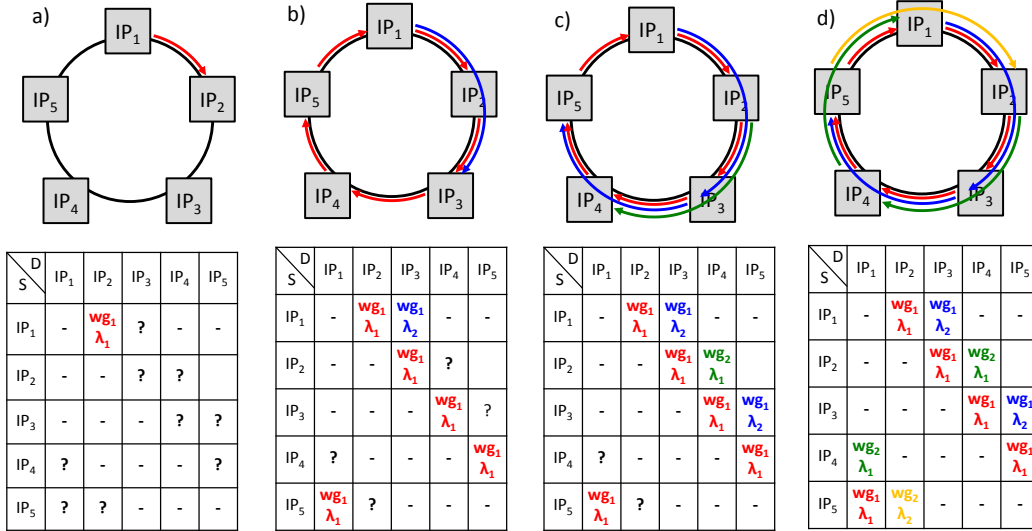


Figure 9: Wavelength assignments for 5 IP cores and for a maximum of two wavelengths per waveguide.

5. MULTI-LAYER IMPLEMENTATIONS OF RELATED WRONOCs

In order to compare ORNoC_{ML} with related WRNoCs, we investigate the design of Matrix [15], λ-router [4] and Snake [16] with multiple optical layers.

5.1 Matrix

Figure 10-a illustrates a multi-layer implementation of Matrix used to interconnect four cores. Waveguide crossings are avoided by allocating inputs and outputs waveguides on the first and the second layer respectively. For its implementation, Matrix uses 16 MRs to fully interconnect the 4 cores. The MRs located on the diagonal can be removed if only inter-core communications are considered, which leads to $(N^2-1) \times N^2$ MRs for $N \times N$ cores architecture.

In order to match with the layout constraints from regular $N \times N$ architecture, Matrix is located in the middle of the optical layer for layout symmetry purposes, as illustrated in Figure 10-b and -c. The ONI transmitter part and receiver part must be connected to the Matrix input and output respectively. Achieving an optimal layout is not an easy task. It depends on system-level parameters (e.g., number of cores and distance between the cores) and technological parameters (e.g., insertion losses). For instance, if $P_{propagation}$ is high (e.g., 2dB/cm), a layout with waveguide crossings but shorter waveguides may show lower total losses L_{total} than a layout without waveguide crossings but with longer waveguides. Therefore, for a fair comparison with ORNoC_{ML}, which avoids waveguide crossings in the same layer, we assumed two layouts. The first layout, shown in Figure 10-b, avoids waveguide crossings and is named Matrix_{ML}^{w/oX}. The second layout, shown in Figure 10-c, minimizes the waveguide length and is named Matrix_{ML}^{wX}.




Figure 10: a) Matrix topology, layouts b) without waveguide crossings and c) with the shortest waveguide length.

5.2 λ -router and Snake

λ -router and Snake are multi-stage optical networks that can be implemented in similar way, as illustrated in Figure 11-a and -b. The optical signals propagate along the waveguides and are dropped from a waveguide to another, in order to reach the targeted outputs. The switching structure of λ -router and Snake is a symmetric PSE implemented with two identical MRs. The

method proposed in [4] is also used: by managing only the required communications, the unnecessary PSEs are removed, which helps reducing the network complexity. By considering only inter-core communications, the PSEs located in the central row and the central column of λ -router and Snake are removed, respectively.

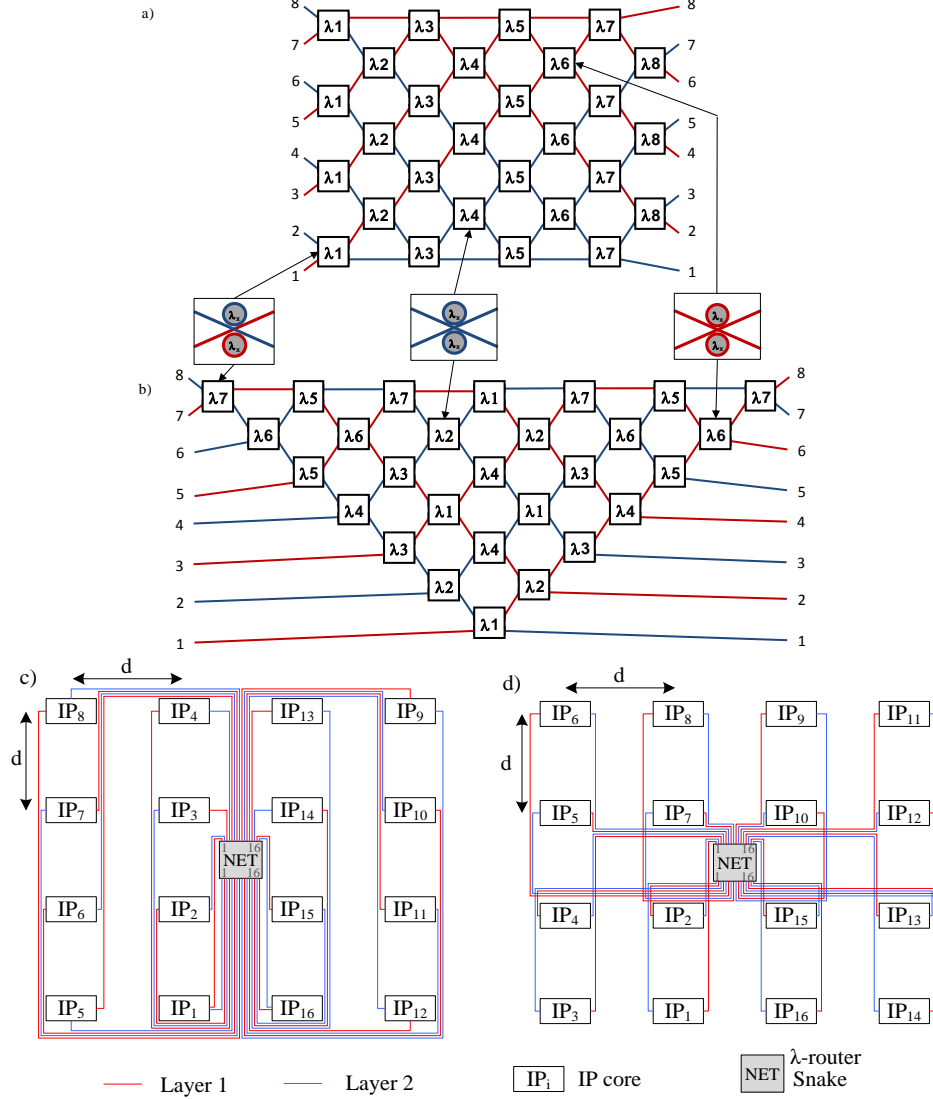


Figure 11: Topology of a) λ -router and b) Snake, and layouts c) without waveguide crossings and d) with the shortest waveguide length.

Multi-stage topologies lead to a significant number of waveguide crossings in the worst-case path. Indeed, for networks with N^2 inputs, there are N^2-1 and $2N^2-5$ waveguide crossings in the worst-case path of λ -router and Snake, respectively. This can be significantly reduced by assuming two-layer implementations illustrated in Figure 11-a and -b. For the sake of regularity and symmetry, the input waveguides are alternately located in the first and second layers. By using this layout design rule, for a 4x4 architecture size of λ -router and Snake crossbars, the

number of waveguide crossings in the worst-case path is reduced from 15 and 27 to 12 and 13, respectively. It represents 20% and 51.9% reduction separately. PSEs with waveguides located in different layers are implemented as described in Figure 3-d.

Similarly to Matrix, the inputs and outputs of the network (located in the center of the optical layer) are connected to the ONIs assuming two layouts. The first layout, shown in Figure 11-c, avoids waveguide crossings and leads to λ -router_{ML}^{w/oX} and Snake_{ML}^{w/oX}. The second layout, shown in Figure 11-d, minimizes the waveguide length and corresponds to λ -router_{ML}^{wX} and Snake_{ML}^{wX}. The layouts will be compared in the result section of the paper.

6. COMPARATIVE STUDY AND RESULTS

We evaluate and compare the multi-layer implementations according to the worst-case loss and average loss metrics. We first discuss the technology related values to be used for the comparisons. In Section 6.2 and 6.3, we compare the best networks (i.e., Matrix_{ML}^{wX} and ORNoC_{ML}) by exploring system-level and technology-level parameters. Also, we evaluate the laser power saving achieved thanks to the multi-layer based implementation of the optical crossbars. Finally, we give a summary of the results and we discuss the results.

6.1 Design parameters

We assume c-Si material only for the implementation of single-layer interconnects (the insertion loss of the first layer from Biberman [27] is assumed for the material). Regarding the multi-layer implementations, we consider the insertion losses parameters from Biberman [27] and Huang [34] (Table 3), by assuming c-Si and Si₃N₄ materials for layer 1 and 2 respectively. For both layers, $P_{crossing}$ =0.05dB and $P_{drop,1}$ =0.5dB. For the multi-layer implementations of λ -router, Snake and Matrix, we evaluate the worst-case and average losses for each communication path following equation (1). This is achieved by evaluating four parameters: i) the signal propagation distance in both layers ($l_{s-d,1}$, $l_{s-d,2}$); ii) the number of waveguide crossing ($N_{crossing}$); iii) the drop operation ($N_{drop,1}$, $N_{drop,2}$); and iv) the inter-layer coupling (N_{ovc}). Regarding ORNoC_{ML}, we follow the design method defined in Section 4.2 for the two sets of parameters. For a 4x4 architecture, the ring assignments obtained for Biberman and Huang parameters are given in Figure 12 -a and -b respectively. In both cases, most of the communications are allocated to layer 2 since it leads to the lowest propagation losses (0.1dB/cm w.r.t. 0.5dB/cm in Figure 12-a; 1.3dB/cm w.r.t. 2.85dB/cm in Figure 12-b). In Figure 12-a, slightly more communications are allocated to layer 1 (17.5%) compared to Figure 12-b (12.5%), which is due to smaller vertical coupling losses. This demonstrates the ability of our design method to assign communications on layer and direction according to technological parameters.

Table 3 Insertion Loss Values

	$P_{propagation,1}$ (dB/cm)	$P_{propagation,2}$ (dB/cm)	P_{ovc} (dB)	$P_{drop,2}$ (dB)
Biberman [27]	0.5	0.1	0.1	0.6
Huang [34]	2.85	1.3	0.2	0.7

a) Biberman [27]

Layer/ Direction	Destination IP core															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Source IP core	1	-	C	C	C	C	C	C	C	C	C	C	C	C	C	C
	2	C	-	C	C	C	C	C	C	C	C	C	C	C	C	C
	3	C	C	-	C	C	C	C	C	C	C	C	C	C	C	C
	4	C	C	C	-	C	C	C	C	C	C	C	C	C	C	C
	5	C	C	C	C	-	C	C	C	C	C	C	C	C	C	C
	6	C	C	C	C	C	-	C	C	C	C	C	C	C	C	C
	7	C	C	C	C	C	C	-	C	C	C	C	C	C	C	C
	8	C	C	C	C	C	C	C	-	C	C	C	C	C	C	C
	9	C	C	C	C	C	C	C	C	-	C	C	C	C	C	C
	10	C	C	C	C	C	C	C	C	C	-	C	C	C	C	C
	11	C	C	C	C	C	C	C	C	C	C	-	C	C	C	C
	12	C	C	C	C	C	C	C	C	C	C	C	-	C	C	C
	13	C	C	C	C	C	C	C	C	C	C	C	C	-	C	C
	14	C	C	C	C	C	C	C	C	C	C	C	C	C	-	C
	15	C	C	C	C	C	C	C	C	C	C	C	C	C	C	-
	16	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C

b) Huang [34]

Layer/ Direction	Destination IP core															
	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15	16
Source IP core	1	-	C	C	C	C	C	C	C	C	C	C	C	C	C	C
	2	C	-	C	C	C	C	C	C	C	C	C	C	C	C	C
	3	C	C	-	C	C	C	C	C	C	C	C	C	C	C	C
	4	C	C	C	-	C	C	C	C	C	C	C	C	C	C	C
	5	C	C	C	C	-	C	C	C	C	C	C	C	C	C	C
	6	C	C	C	C	C	-	C	C	C	C	C	C	C	C	C
	7	C	C	C	C	C	C	-	C	C	C	C	C	C	C	C
	8	C	C	C	C	C	C	C	-	C	C	C	C	C	C	C
	9	C	C	C	C	C	C	C	C	-	C	C	C	C	C	C
	10	C	C	C	C	C	C	C	C	C	-	C	C	C	C	C
	11	C	C	C	C	C	C	C	C	C	C	-	C	C	C	C
	12	C	C	C	C	C	C	C	C	C	C	C	-	C	C	C
	13	C	C	C	C	C	C	C	C	C	C	C	C	-	C	C
	14	C	C	C	C	C	C	C	C	C	C	C	C	C	-	C
	15	C	C	C	C	C	C	C	C	C	C	C	C	C	C	-
	16	C	C	C	C	C	C	C	C	C	C	C	C	C	C	C

Figure 12: Ring assignment in ORNoC_{ML} by assuming losses values from a) Biberman [27] and b) Huang [34].

6.2 Crossbar comparison under system-level parameters exploration

6.2.1 Architecture sizes

We assume a fixed 2cm×2cm die size as in [21] and we evaluate the losses for 2×2, 4×4, 6×6 and 8×8 architecture sizes, i.e., distance between neighboring IP cores $d=10, 5, 3.33$ and 2.5mm, respectively. All the results of this section are given for technological parameters from Biberman [27] (listed in Table 3).

In Figure 13, we first estimate the worst-case and average loss reductions (in %) for the two-layer implementation over the single-layer implementation for Matrix_{ML}^{w/oX}, Matrix_{ML}^{wX}, λ -router_{ML}^{w/oX}, λ -router_{ML}^{wX}, Snake_{ML}^{w/oX}, Snake_{ML}^{wX}, and ORNoC_{ML} crossbars. 0% means that multi-layer and single-layer implementations lead to the same losses. Results above 0% indicate a reduction of the losses for the multi-layer implementation. Figure 13-a shows that improvements are obtained even for the smallest size architecture: the reduction of waveguide crossings allows compensating the vertical coupling losses. For instance, slight reduction of worst-case losses is obtained for 2×2 architecture size: Matrix_{ML}^{w/oX} (31%), Matrix_{ML}^{wX} (24%), λ -router_{ML}^{w/oX} (4%), λ -router_{ML}^{wX} (11%), Snake_{ML}^{w/oX} (4%), Snake_{ML}^{wX} (11%), and ORNoC_{ML} (40%). For 8×8 size, the improvements of Matrix_{ML}^{wX}, λ -router_{ML}^{wX} and Snake_{ML}^{wX} reach 69%, 28% and 42%, respectively. The improvement for Snake is higher than for λ -router due to the initially higher number of waveguide crossings. Overall, Matrix demonstrates better improvement compared to λ -router and Snake since there is no additional waveguide crossings. The layout with the shortest waveguide length shows better improvement since it directly takes benefits from the reduction of the number of waveguide crossings. In the meanwhile, ORNoC_{ML} achieves a 67% improvement since the propagation loss in layer 2 is much lower than in layer 1 (Table 3).

A similar trend is observed for the average loss (Figure 13-b). Matrix shows the largest improvement among Matrix, λ -router and Snake, since its single-layer implementation exhibits the highest number of waveguide crossings. For example, for 8×8 Matrix_{ML}^{wX}, the two-layer implementation allows to reduce the number of waveguide crossings from 125 to 88. ORNoC also demonstrates significant improvement due to the lower propagation loss in layer 2. As an example, the reductions in the worst-case and average losses reach 67% and 58% respectively for 8×8 IP cores.

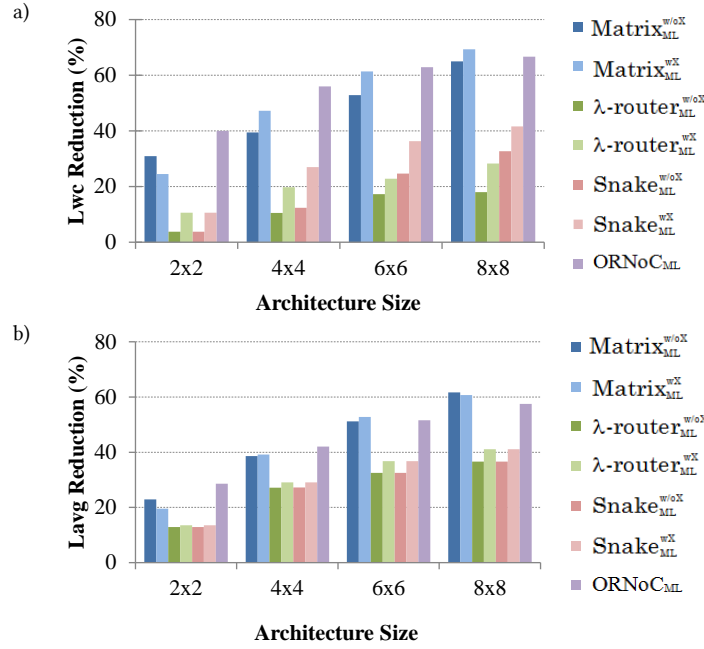


Figure 13: Improvement of multi-layer implementations of Matrix, λ-router, Snake and ORNoC against the single-layer implementations considering: a) worst-case losses and b) average losses.

Figure 14-a and -b detail the loss contribution to the worst-case and average loss respectively for Matrix_{ML}^{w/oX}, Matrix_{ML}^{wX}, λ-router_{ML}^{w/oX}, λ-router_{ML}^{wX}, Snake_{ML}^{w/oX}, Snake_{ML}^{wX}, and ORNoC_{ML}. A first observation on the worst-case loss evaluation can be made regarding the layouts: for 2×2, 4×4 and 6×6 architecture size, the layout with the shortest waveguide lengths outperforms the layout without any waveguide crossing, independently from the network topology. However, the layout without any waveguide crossing shows better scalability since the loss shows lower sensibility to the architecture size variation. For 8×8 architecture size, it exhibits lower losses for Matrix, λ-router and Snake. Similar observation can be made for average loss (Figure 14-b).

The results indicate that the better scalability would combine the use of: i) multi-layer deposited silicon technology, to reduce waveguide crossings in the network by implementing multiple optical layers; and ii) intra-layer layout that avoids waveguide crossings. ORNoC_{ML} gathers these criteria, leading to the lowest worst-case loss despite the long distance introduced by the serpentine layout. For the 8×8 case, the worst-case path in ORNoC_{ML} is 1.5dB, lower than Matrix_{ML}^{w/oX} with 3.3dB and Matrix_{ML}^{wX} with 3.7dB.

By considering the average loss, ORNoC_{ML} reduces the average loss by 63% on average compared to the other multi-layer implementations. This significant difference is obtained due to the shorter propagation distance between neighbor IP cores. The improvement reaches 55% and 70% for 2×2 case and 8×8 case, respectively. The average loss is 1.1dB for ORNoC_{ML} compared to 2.4dB and 3.2dB for Matrix_{ML}^{w/oX} and Matrix_{ML}^{wX} under the 8×8 architecture size.

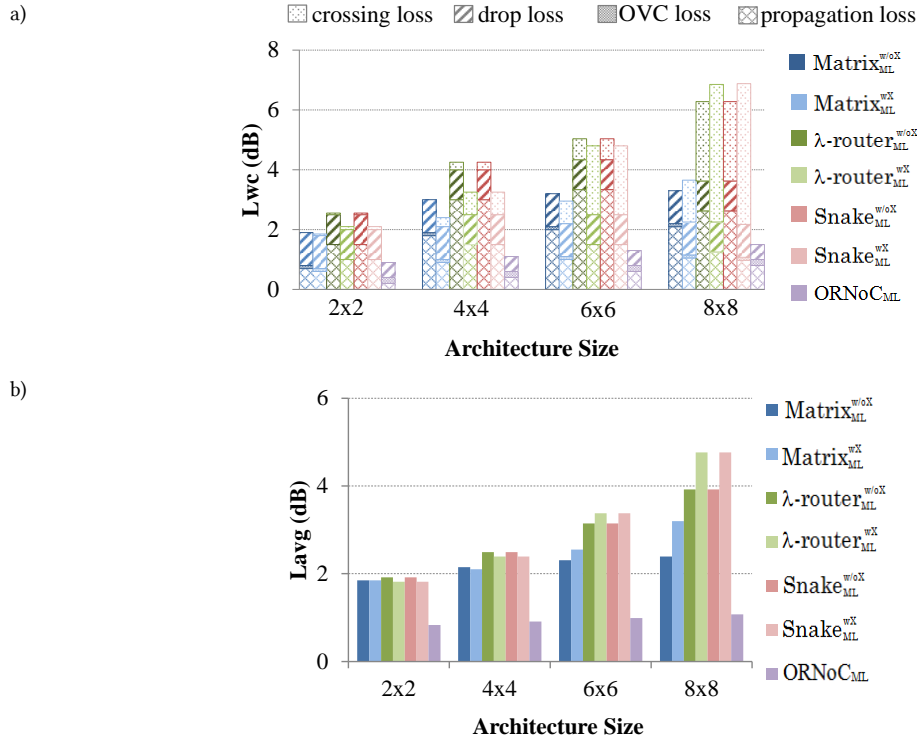


Figure 14: a) Worst-case losses and b) average losses evaluation for 2x2 to 8x8 IP cores.

6.2.2 Distance between the cores

Figure 15-a shows the comparison results for a fixed 6×6 cores with d ranging from 1mm to 3mm, with intervals of 0.5mm. The increase of the loss with the distance is higher for the networks relying on the layout without any waveguide crossing. In all the cases, even for the longest considered distance (i.e., 3mm, which leads to a 3.24cm² die size), ORNoC_{ML} is the most power-efficient network and is followed by Matrix_{ML}^{wX} and Matrix_{ML}^{w/oX}. Similar trend is observed for the average loss in Figure 15-b.

For the 8×8 size, the implementation of Matrix requires 63 wavelengths with regard to 64 wavelengths for Snake and λ-router. Architectures which include higher number of wavelengths are penalized by the crosstalk and fabrication variability. A more reasonable implementation would be to consider several smaller networks, which implies additional waveguide crossings [40]. The use of the ring topology intrinsically leverages this issue since the number of waveguides can be set according to the crosstalk and process variability requirements. This can be achieved without any waveguide crossing, because of the multi-layer implementation and the use of on-chip laser sources.

Following the methodology from [5], ORNoC_{ML} would require 16 waveguides if we consider the optimistic maximum number of 64 wavelengths per waveguides, and 63 waveguides if we consider more realistic scenario with 16 wavelengths per waveguide. When parallel waveguides are added for Matrix, λ-router or Snake, additional waveguide crossings are introduced [40], even when multi-layer technology is employed. For the ORNoC_{ML}, no additional waveguide crossing is included. This characteristic together with the regularity of its layout turns ORNoC_{ML} into a scalable structure which does not require any custom place-and-route tool [16][41].




Figure 15: Crossbar comparison for 6x6 cores according to a) worst-case losses and b) average losses with distance between cores ranging from 1mm to 3mm.

6.2.3 Laser output power saving

The minimum laser output power required for the communication is evaluated for the two most energy-efficient architectures, i.e., ORNoC_{ML} and $\text{Matrix}_{\text{ML}}^{\text{wX}}$. We assume 80% laser coupling efficiency. The laser output power saving ratio for ORNoC_{ML} over $\text{Matrix}_{\text{ML}}^{\text{wX}}$ is shown in Figure 16. For instance, for 2x2 architecture and 2.5mm between the cores, the required laser output power for ORNoC_{ML} is reduced by 14% compared to the solution with $\text{Matrix}_{\text{ML}}^{\text{wX}}$. Results show that power saving under a given architecture size remains similar. However, significant saving is achieved for larger architectures: for a 2.5mm distance, the laser power saving increases from 15% (2x2) to 37% (8x8). The improvement is due to increasing number of waveguide crossings with $\text{Matrix}_{\text{ML}}^{\text{wX}}$. It is worth noticing that, these results being provided for the average losses in the communications path, additional power saving could be achieved for ORNoC_{ML} if tunable lasers output power are used [31].




Figure 16: Laser output power saving for ORNoC_{ML} over Matrix_{ML}^{wX}.

6.3 Crossbar comparisons under technological parameters exploration

Comparisons achieved in the previous sections are based on a given set of losses values. Such analysis may lead to incomplete and/or unfair comparisons. For instance, by considering low propagation losses and high waveguide crossing losses values, layouts without any waveguide crossing will be favored over the layout with the shortest waveguide length. For this purpose, we further compare ORNoC_{ML} and Matrix_{ML}^{wX} (i.e., the best networks based on the previous analysis) by exploring technology related parameters. The following results are given for the worst-case loss evaluation under 8x8 architecture sizes.

6.3.1 Exploration through propagation loss and OVC loss parameters

For the first comparison, $P_{propagation,1}$ and $P_{propagation,2}$ are ranged from 0 to 3dB/cm and from 0 to 1.5dB/cm, respectively. Figure 17 show the worst-case loss for Matrix_{ML}^{wX} (blue color) and ORNoC_{ML} (green color) assuming 1mm, 1.5mm, 2mm and 2.5mm distances between IP cores. For instance, for $d=1$ mm (Figure 17-a), $P_{propagation,1}=0.5$ dB/cm and $P_{propagation,2}=0.1$ dB/cm, worst-case losses for ORNoC_{ML} and Matrix_{ML}^{wX} are 1.0dB and 3.0dB respectively. The worst-case loss for Matrix_{ML}^{wX} increases linearly with the propagation loss. The trend is different for ORNoC_{ML} for which communications are allocated on the path showing lower losses: for $P_{propagation,2}$ smaller than $P_{propagation,1}$, layer 2 is utilized in priority.

For $d=1$ mm, ORNoC_{ML} outperforms Matrix_{ML}^{wX} for most propagations losses values, including those extracted from [27] and [34]. However, Matrix_{ML}^{wX} shows lower losses than ORNoC_{ML} (4.7dB and 5.3dB respectively) around $P_{propagation,1}=P_{propagation,2}=1.5$ dB values. Obviously, worst-case losses for both Matrix_{ML}^{wX} and ORNoC_{ML} tend to increase with larger distance between IPs, which is due to the increased waveguides length. However, the serpentine layout of the ring topology is more impacted by the increased lengths and Matrix_{ML}^{wX} becomes more efficient than ORNoC_{ML}. However, this trend is limited to region for which the ratio between the propagation losses on the layers remains small. This is further investigated in the following.




Figure 17: Exploration of $\text{Matrix}_{\text{ML}}^{\text{wX}}$ (in blue) and ORNoC_{ML} (in green) worst-case loss according to propagation loss parameters for 8×8 IP cores and four distances: a) $d=1\text{mm}$, b) $d=1.5\text{mm}$, c) $d=2\text{mm}$ and d) $d=2.5\text{mm}$. Results are given for $P_{\text{OVC}}=0.1\text{dB}$, $P_{\text{drop},i}=0.5\text{dB}$, and $P_{\text{crossing}}=0.05\text{dB}$.

The intersection lines from Figure 17 (i.e., when worst-case losses of $\text{Matrix}_{\text{ML}}^{\text{wX}}$ and ORNoC_{ML} are the same) are reported in Figure 18-a. In the figure, each line corresponds to a distance (i.e., $d=1\text{mm}$, 1.5mm , 2mm , and 2.5mm). The left-hand side of a line is the area for which ORNoC_{ML} is more energy efficient than $\text{Matrix}_{\text{ML}}^{\text{wX}}$. For Huang [34] propagation loss parameters, $\text{Matrix}_{\text{ML}}^{\text{wX}}$ is more energy efficient than ORNoC_{ML} for $d=2\text{mm}$ and $d=2.5\text{mm}$ while, for much lower losses parameters from Biberman [27], ORNoC_{ML} dominates over $\text{Matrix}_{\text{ML}}^{\text{wX}}$ for all the distances.

The comparison for 6×6 architecture is illustrated in Figure 18-b. While the trend is similar to the one obtained for 8×8 , ORNoC_{ML} is more energy efficient than $\text{Matrix}_{\text{ML}}^{\text{wX}}$ for most design options. As a result, for Huang [34] values, ORNoC_{ML} is the most energy efficient solution, independently from the distance. The intersection line for $d=1\text{mm}$ is out of the studied propagation loss ranges. This trend is compatible with the observation made in Section 6.2 and can be summarized as follow: the shift from 8×8 to 6×6 architecture size leads to i) a reduction in the waveguide crossing for $\text{Matrix}_{\text{ML}}^{\text{wX}}$ and ii) reduced waveguide length for ORNoC_{ML} . The design of ORNoC_{ML} being optimized according to the propagation losses, significant improvements are obtained compared to a naïve allocation of the communication. However, the energy improvement compared to a network with waveguide crossings depends on the crossing losses, which is investigated in the following.

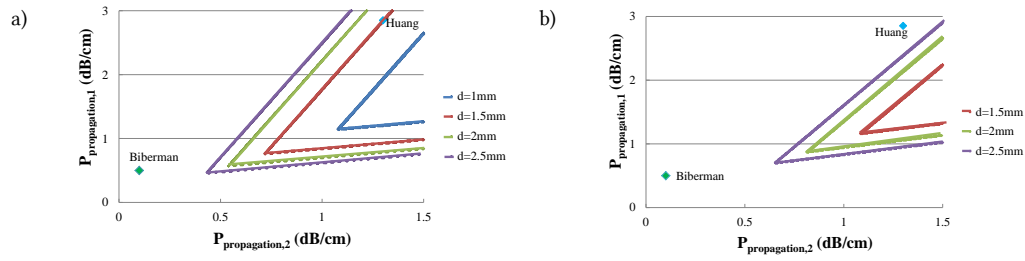


Figure 18: Energy efficiency comparison of $\text{Matrix}_{\text{ML}}^{\text{wX}}$ and ORNoC_{ML} according to the propagation loss in layer 1 and layer 2 for a) 8×8 and b) 6×6 cores. The lines represent the values for which $\text{Matrix}_{\text{ML}}^{\text{wX}}$ and ORNoC_{ML} show the same worst-case loss for $d=1\text{mm}$, 1.5mm , 2mm and 2.5mm . For each distance, ORNoC_{ML} is more energy efficient on the *left*-hand side of the line. Results are given for 8×8 IP cores, $P_{\text{drop},1}=0.5\text{dB}$, $P_{\text{crossing}}=0.05\text{dB}$, $P_{\text{OVC}}=0.1\text{dB}$.

6.3.2 Comparison through propagation loss and crossing losses parameters




Figure 19: Energy efficiency comparison of $\text{Matrix}_{\text{ML}}^{\text{wX}}$ and ORNoC_{ML} according to the crossing loss and the propagation loss ratio. The lines represent the values for which $\text{Matrix}_{\text{ML}}^{\text{wX}}$ and ORNoC_{ML} show the same worst-case loss for $d=1\text{mm}$, 1.5mm , 2mm and 2.5mm . For each distance, ORNoC_{ML} is more energy efficient on the *right*-hand side of the line. Results are given for 8×8 IP cores, $P_{\text{OVC}}=0.2\text{dB}$, $P_{\text{drop},1}=0.5\text{dB}$, $P_{\text{propagation},2}=1.3\text{dB/cm}$.

In the following, we further investigate the comparison between the two interconnects by exploring the crossing loss (i.e., P_{crossing} in the $0\text{--}0.2\text{dB}$ range [6]) and the propagation loss ratio among the layers (i.e., $P_{\text{propagation},1}/P_{\text{propagation},2}$). $P_{\text{propagation},2}$ is set to 1.3dB/cm [34] and we assume lower losses in layer 2 than in layer 1. Figure 19 illustrates the results for 8×8 cores. We use the same representation as in Figure 18. The area on the right of the line corresponds to design space for which ORNoC_{ML} is more energy efficient than $\text{Matrix}_{\text{ML}}^{\text{wX}}$. The results help to understand the impact of the crossing and the propagation losses on the network energy efficiencies. For instance, for $d=2.5\text{mm}$ and 0.15dB crossing loss, ORNoC_{ML} is more energy efficient than $\text{Matrix}_{\text{ML}}^{\text{wX}}$ from a 1.6 propagation ratio. However, if the crossing loss can be reduced to 0.05dB without modification of the propagation loss ratio, then $\text{Matrix}_{\text{ML}}^{\text{wX}}$ is the best network and ORNoC_{ML} should be used only if a 2.8 ratio can be reached.

6.4 Summary of the main results and discussion

The results have shown that multi-layer deposited silicon technology contributes to improve the energy efficiency of optical crossbars thanks to drastic losses reduction. The most significant loss reductions have been observed for layouts that minimize the waveguide length while still

allowing waveguide crossings (e.g., $\text{Matrix}_{\text{ML}}^{\text{wX}}$, $\lambda\text{-router}_{\text{ML}}^{\text{wX}}$ and $\text{Snake}_{\text{ML}}^{\text{wX}}$). Furthermore, the bigger the architecture, the more the energy saving, e.g., 70% worst-case loss reduction is reached for $\text{Matrix}_{\text{ML}}^{\text{wX}}$ interconnecting 8×8 cores.

We compared all the multi-layer implementations and by aggregating the contribution of the propagation loss, OVC loss, drop loss and crossing loss in the worst-case loss. Overall, ORNoC_{ML} provides the lowest worst-case loss for 2×2 to 8×8 architectures, overcoming the multi-layer implementations of Matrix, λ -router and Snake. The higher energy efficiency of ORNoC_{ML} is due to i) the serpentine layout (that avoids waveguide crossings) and ii) the design method (that allows allocating communication on optical path showing the lowest losses). For instance, ORNoC_{ML} achieves on average 55% and 60% reduction of worst-case and average losses, compared to $\text{Matrix}_{\text{ML}}^{\text{wX}}$. For larger architecture size, e.g., 10×10 , Matrix, Snake and λ -router will reach a physical limitation related to the maximum number of wavelengths per waveguide (e.g., 64 wavelengths [42]). A solution to overcome this limitation is to replicate the network implementation, which leads to additional waveguide crossings and less regular layout. Using additional optical layers (not only two as in this study) could also help to overcome this issue.

We also investigate the energy efficiency comparison of $\text{Matrix}_{\text{ML}}^{\text{wX}}$ and ORNoC_{ML} under technological parameters exploration. The results can be used in two ways: first, from a set of technological parameters, it is possible to identify the best topology; second, by targeting a given topology, constraints on technological parameters can be identified.

7. CONCLUSION

In this paper, we have investigated the impact of multi-layer deposited silicon technology on the energy efficiency of wavelength-routed optical network-on-chip. For the ring topology, a design method has been proposed; it allows allocating communications on optical paths showing the lowest losses. For Matrix, λ -router and Snake topologies, we proposed layouts i) with minimized waveguide length and ii) without waveguide crossing. Results show that, to interconnect 8×8 cores, multi-layer implementations lead to on average 42% and 46% reduction in the worst-case and averages losses, respectively. This has an immediate impact on the laser output power, which can be decreased to up to 85%, thus contributing to the higher energy efficiency of the optical network. The ring is the most energy-efficient topology among all studied architectures: on average, it leads to 66% reduction of worst-case loss when compared to the related topologies. We also investigated the impact of technological parameters values on ORNoC and Matrix energy efficiency. This allows selecting the topology to be used for a given technological platform. In our future work, we will further investigate the impact of multi-layer silicon deposited on the network thermal sensitivity and robustness to fabrication process variation.

ACKNOWLEDGMENT

Hui LI is supported by China Scholarship Council (CSC). This work has received a French Government support granted to the COMIN Labs excellence laboratory and managed by the National Research Agency in the "Investing for the Future" program under reference ANR-10-LABX-07-01.

REFERENCES

- [1] C. Sun, M. T. Wade, Y. Lee, J. S. Orcutt, L. Alloatti, M. S. Georgas, A. S. Waterman, J. M. Shainline, R. R. Avizienis, S. Lin, B. R. Moss, R. Kumar, F. Pavanello, A. H. Atabaki, H. M. Cook, A. J. Ou, J. C. Leu, Y.-H. Chen, K. Asanović, R. J. Ram, M. a. Popović, and V. M. Stojanović, "Single-chip microprocessor that communicates directly using light," *Nature* 528, pp. 534–538, 2015.

- [2] J. Psota, J. Miller, G. Kurian, H. Hoffman, N. Beckmann, J. Eastep, and A. Agarwal, "ATAC: Improving Performance and Programmability With on-Chip Optical Networks," In Proceedings of IEEE International Symposium on Circuits and Systems (ISCAS), pages 3325–3328, 2010.
- [3] A. Shacham, K. Bergman and L. P. Carloni, "Photonic Networks-on-Chip for Future Generations of Chip Multiprocessors," in IEEE Transactions on Computers, Vol. 57, No. 9, pp. 1246–1260, 2008.
- [4] I. O'Connor, F. Mieyeville, F. Gaffiot, A. Scandurra, and G. Niolescu, "Reduction Methods for Adapting Optical Network on Chip Topologies to Specific Routing Applications," In Proceedings of DCIS, 2008.
- [5] S. Le Beux, J. Trajkovic, I. O'Connor and G. Niolescu, "Layout guidelines for 3D architectures including Optical Ring Network-on-Chip (ORNoC)," in 2011 IEEE/IFIP 19th International Conference on VLSI and System-on-Chip, pp. 242–247, 2011.
- [6] P. Koka, M. O. McCracken, H. Schwetman, C.-H. O. Chen, X. Zheng, R. Ho, K. Raj, and A. V. Krishnamoorthy, "A micro-architectural analysis of switched photonic multi-chip interconnects," In 39th Annual International Symposium on Computer Architecture, 2012.
- [7] A. Biberman, N. Sherwood-Droz, X. Zhu, M. Lipson, and K. Bergman, "High-Speed Data Transmission in Multi-Layer Deposited Silicon Photonics for Advanced Photonic Networks-on-Chip," in CLEO:2011 - Laser Applications to Photonic Applications, OSA Technical Digest (CD) (Optical Society of America, 2011), paper CThA1.
- [8] R. Hendry, G. Hendry, and K. Bergman, "TDM Photonic Network Using Deposited Materials," High Performance Embedded Computing (HPEC), 2011.
- [9] R. Sun, M. Beals, A. Pomerene, J. Cheng, C.-y. Hong, L. Kimerling, and J. Michel, "Impedance matching vertical optical waveguide couplers for dense high index contrast circuits," Optics Express, Vol. 16, No. 16, pp. 11682–11690, 2008.
- [10] A. Parini, G. Calò, G. Bellanca, and V. Petruzzelli, "Vertical link solutions for multilayer optical-networks-on-chip topologies," Optical and Quantum Electronics, Vol. 46, Issue 3, pp. 385–396, 2014.
- [11] N. Sherwood-Droz, and M. Lipson, "Scalable 3D dense integration of photonics on bulk silicon," Optics Express, Vol. 19, No. 18, 2011.
- [12] J. T. Bessette and D. Ahn, "Vertically stacked microring waveguides for coupling between multiple photonic planes," Opt. Express, Vol. 21, No.11, pp. 13580–13591, 2013.
- [13] G. Calò and V. Petruzzelli, "Wavelength routers for multilayer integrated optical networks on chip," in proceedings of 2015 17th International Conference on Transparent Optical Networks (ICTON), 2015.
- [14] G. Calò and V. Petruzzelli, "Generic Wavelength-routed Optical Router (GWOR) based on grating-assisted vertical couplers for multilayer optical networks," Optics Communications, Vol. 366, pp. 99–106, 2016.
- [15] A. Bianco, D. Cuda, M. Garrich, R. Gaudino, G. Gavilanes, P. Giaccone, and F. Neri, "Optical Interconnection Networks based on Microring Resonators," In Proceedings of IEEE International Conference on Communications, 2010.
- [16] L. Ramini, P. Grani, S. Bartolini, and D. Bertozzi, "Contrasting wavelength-routed optical NoC topologies for power-efficient 3d-stacked multicore processors using physical-layer analysis," in Proceedings of Design, Automation & Test in Europe Conference & Exhibition (DATE), pp. 1589–1594, 2013.
- [17] S. Pasricha and S. Bahirat, "OPAL: A multi-layer hybrid photonic NoC for 3D ICs," 16th Asia and South Pacific Design Automation Conference (ASP-DAC 2011), 2011.
- [18] D. Dang, B. Patra and R. Mahapatra, "A 2-layer laser multiplexed photonic network-on-chip," Sixteenth International Symposium on Quality Electronic Design, 2015.
- [19] K. Chen, H. Gu, Y. Yang, and D. Fan, "A Novel Two-Layer Passive Optical Interconnection Network for On-Chip Communication," Journal of Lightwave Technology, Vol. 32, No. 5, 2014.
- [20] S. Le Beux, H. Li, G. Niolescu, J. Trajkovic, and I. O'Connor, "Optical crossbars on chip, a comparative study based on worst-case losses," Concurrency and Computation: Practice and Experience, Vol. 26, No. 15, pp. 2492–2503, 2014.
- [21] X. Zhang and A. Louri. "A Multilayer Nanophotonic Interconnection Network for On-Chip Many-core Communication," in Proceedings for DAC, 2010.
- [22] R. Morris, A. K. Kodi, and A. Louri. "Dynamic Reconfiguration of 3D Photonic Networks-on-Chip for Maximizing Performance and Improving Fault Tolerance," in IEEE/ACM 45th Annual International Symposium on Microarchitecture, 2012.
- [23] R. W. Morris, A. K. Kodi, A. Louri and R. D. Whaley, "Three-Dimensional Stacked Nanophotonic Network-on-Chip Architecture with Minimal Reconfiguration," in IEEE Transactions on Computers, Vol. 63, No. 1, pp. 243–255, 2014.
- [24] Y. Pan, J. Kim and G. Memik, "FlexiShare: Channel sharing for an energy-efficient nanophotonic crossbar," in The Sixteenth International Symposium on High-Performance Computer Architecture (HPCA – 16), 2010.
- [25] Z. Chen, H. Gu, Y. Yang, and D. Fan, "A Hierarchical Optical Network-On-Chip Using Central-Controlled Subnet and Wavelength Assignment," Journal of Lightwave Technology, Vol. 32, No. 5, 2014.
- [26] H. Li, S. Le Beux, G. Niolescu, and I. O'Connor. "Energy-Efficient Optical Crossbars on Chip with Multi-Layer Deposited Silicon," In Proceedings of ASP-DAC, 2015.
- [27] A. Biberman, K. Preston, G. Hendry, N. Sherwood-Droz, J. Chan, J. S. Levy, M. Lipson, and K. Bergman, "Photonic Network-on-Chip Architectures Using Multilayer Deposited Silicon Materials for High-Performance Chip Multiprocessors," ACM Journal on Emerging Technologies in Computing Systems, Vol. 7, No. 2, pp. 7:1–7:25, 2011.

- [28] I. Loi, F. Angiolini, and L. Benini, "Supporting Vertical Links for 3D Networks-on-Chip: Toward an Automated Design and Analysis Flow," In Proceedings of the 2nd international conference on Nano-Networks (Nano-Net'07), pages 1–5, 2007.
- [29] J. V. Campenhout, L. Liu, P. R. Romeo, D. V. Thourhout, C. Seassal, P. Regreny, L. D. Cioccio, J.-M. Fedeli, and R. Baets, "A compact SOI-integrated multiwavelength laser source based on cascaded InP microdisks," IEEE Photon. Technol. Lett., Vol. 20, No. 16, pp. 1345–1347, 2008.
- [30] 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 Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA), pages 153–164, 2008.
- [31] H. Li, A. Fourmigue, S. Le Beux, I. O'Connor, and G. Nicolescu, "Towards Maximum Energy Efficiency in Nanophotonic Interconnects with Thermal-Aware On-Chip Laser Tuning," in IEEE Transactions on Emerging Topics in Computing (in publication), 2016.
- [32] A. M. Jones, C. T. DeRose, A. L. Lentine, D. C. Trotter, A. L. Starbuck, and R. A. Norwood, "Ultra-low crosstalk, CMOS compatible waveguide crossings for densely integrated photonic interconnection networks," Opt. Express, Vol. 21, No.10, pp. 12002-12013, 2013.
- [33] R. Schuster, A. Parini, and G. Bellanca, "Parametric exploration of vertical tapered coupler for 3D optical interconnection," in OPTICS workshop, 2015.
- [34] Y. Huang, J. Song, X. Luo, T.-Y. Liow, and G.-Q. Lo, "CMOS compatible monolithic multi-layer Si₃N₄-on-SOI platform for low-loss high performance silicon photonics dense integration," Optics Express, Vol. 22, No. 18, 2014.
- [35] W. D. Sacher, Y. Huang, G. Q. Lo and J. K. S. Poon, "Multilayer Silicon Nitride-on-Silicon Integrated Photonic Platforms and Devices," in Journal of Lightwave Technology, Vol. 33, No. 4, pp. 901-910, 2015.
- [36] V. Donzella, S. T. Fard and L. Chrostowski. "Study of waveguide crosstalk in silicon photonics integrated circuits," in Proc. of SPIE 8915, Photonics North 2013, 89150Z.
- [37] S. Le Beux, J. Trajkovic, I. O'Connor, G. Nicolescu, G. Bois, and P. Paulin, "Optical Ring Network-on-Chip (ORNoC): Architecture and design methodology," in Proceedings of Design, Automation & Test in Europe (DATE'11), 2011.
- [38] C. Sciancalepore, B. B. Bakir, C. Seassal, X. Letartre, J. Harduin, N. Olivier, J.-M. Fedeli, and P. Viktorovitch, "Thermal, Modal, and Polarization Features of Double Photonic Crystal Vertical-Cavity Surface-Emitting Lasers," IEEE Photonics journal, Vol. 4, No 2, pp. 399-410, 2012.
- [39] K. Ohira, K. Kobayashi, N. Iizuka, H. Yoshida, M. Ezaki, H. Uemura, A. Kojima, K. Nakamura, H. Furuyama, and H. Shibata, "On-chip optical interconnection by using integrated III-V laser diode and photodetector with silicon waveguide," Optics Express, Vol. 18, No.15, pp. 15440-15447, 2010.
- [40] S. Le Beux, J. Trajkovic, I. O'Connor, G. Nicolescu, G. Bois, and P. Paulin, "Multi-Optical Network-on-Chip for Large Scale MPSoC," in IEEE Embedded Systems Letters, Vol. 2, No. 3, pp. 77-80, 2010.
- [41] L. Ramini, D. Bertozzi, and L. P. Carloni, "Engineering a Bandwidth-Scalable Optical Layer for a 3D Multi-core Processor with Awareness of Layout Constraints," in 2012 Sixth IEEE/ACM International Symposium on Networks on Chip (NoCS), 2012.
- [42] C. Batten, A. Joshi, J. Orcutt, A. Khilo, B. Moss, C. Holzwarth, M. Popovic, H. Li, H. Smith, J. Hoyt, F. Kartner, R. Ram, V. Stojanovic, and K. Asanovic, "Building Manycore Processor-to-DRAM Networks with Monolithic Silicon Photonics," In HOTI' 08, pages 21-30, 2008.

Received June 2016; revised November 2016 and February 2017;