



HAL
open science

Towards Maximum Energy Efficiency in Nanophotonic Interconnects with Thermal-Aware On-Chip Laser Tuning

Hui Li, Alain Fourmigue, Sébastien Le Beux, Ian O'Connor, Gabriela Nicolescu

► **To cite this version:**

Hui Li, Alain Fourmigue, Sébastien Le Beux, Ian O'Connor, Gabriela Nicolescu. Towards Maximum Energy Efficiency in Nanophotonic Interconnects with Thermal-Aware On-Chip Laser Tuning. IEEE Transactions on Emerging Topics in Computing, 2016, 6 (3), pp.343-356. <10.1109/TETC.2016.2561623>. <hal-01330470>

HAL Id: hal-01330470

<https://inria.hal.science/hal-01330470v1>

Submitted on 10 Jun 2016

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.



HAL Authorization

Towards Maximum Energy Efficiency in Nanophotonic Interconnects with Thermal-Aware On-Chip Laser Tuning

Hui Li¹, Alain Fourmigue², Sébastien Le Beux^{1*}, Ian O'Connor¹ and Gabriela Nicolescu²

¹ Lyon Institute of Nanotechnology, INL-UMR5270

Ecole Centrale de Lyon, Ecully, F-69134, France

² Computer and Software Engineering Dept.

Polytechnique Montréal, Montréal (QC), Canada

* Contact author: sebastien.le-beux@ec-lyon.fr

Abstract—Nanophotonic is an emerging technology considered as one of the key solutions for future generation on-chip interconnects. However, silicon photonic devices are highly sensitive to temperature variation. Under a given chip activity, this leads to a lower laser efficiency and a drift of wavelengths of optical devices (on-chip lasers and microring resonators (MRs)), which results in a higher Bit Error Ratio (BER). In this paper, we propose to jointly tune the on-chip lasers and MRs in order to align the wavelengths of the emitted signals with the resonant wavelengths of the MRs. Our method allows significant improvements of the power consumption with regard to the related methods, while meeting the BER requirement. Compared to methods for which laser tuning is not possible, results show that a combined tuning of laser and MRs leads to 53% energy reduction when the uniform chip activity decreases from 20% to 5%. BER-energy tradeoffs have been explored and allow strategies to be defined to minimize either the energy, or the BER. As a key result, we have shown that, under specific chip activities, increasing the laser power consumption allows both energy and BER to be improved. This trend has been observed for a MWSR channel interconnecting 12 interfaces.

Index Terms— Nanophotonic interconnects, 3D integrated circuits, thermal simulation.



I. INTRODUCTION

TECHNOLOGY scaling down to the ultra-deep submicron domain enables the integration of hundreds of cores. In order to enhance connecting many cores, chip-level communication needs a new kind of interconnect to bring down the power budget. On-chip nanophotonic interconnects are an emerging technology considered as one of the key solutions for future generations of many cores [1], providing high bandwidth and low latency. Interconnect architectures, e.g. MWSR (Multiple Writer Single Reader) [2], are composed of laser sources, microring resonators (MRs), waveguide and photodetector.

Laser sources can be fabricated in several ways and rely, for instance, on Fabry-Perot (FP), distributed feed-back (DFB) and distributed Bragg reflector (DBR) [3]. Among the available technologies, vertical-cavity surface-emitting lasers (VCSELs) provide an optical output power in the range of hundreds of microwatts, thus meeting the requirements for nanophotonic interconnects [3]. Furthermore, VCSELs based on a double set of Si/SiO₂ photonic crystal mirrors (PCMs), so called PCM-VCSELs, are CMOS-compatible [4][5]. They can be used as on-chip laser sources since their footprint is of the same order of magnitude as the size of a MR used to modulate continuous waves emitted by off-chip lasers. VCSELs are thus sufficiently compact to be implemented in a large number. In this way, it saves the laser power since losses caused by power waveguides [2] are avoided, compared with off-chip lasers.

However, silicon photonic devices are highly sensitive to temperature variation induced by thermal effect over

the chip, which leads to a drift of the laser wavelength and MRs resonant wavelength along a same photonic channel. Consequently, the Signal to Noise Ratio (SNR) of the signal received by the photodetector decreases, which leads to a higher Bit Error Ratio (BER). This is further accentuated by the significant reduction of on-chip lasers efficiency as the temperature increases.

This paper proposes a thermal-aware tuning method to overcome the wavelength variation induced by temperature variation, while achieving the targeted BER. The novelty of the method relies on the tuning of the laser driver current, which ideally complements traditional methods such as MRs tuning and channel remapping [14][15]. While the method is evaluated for on-chip laser sources, it is also suitable for off-chip lasers since the impact of a temperature elevation would remain the same.

The paper is structured as follows. Section II presents the related work. Section III describes the problem induced by thermal variation and our proposed method. Section IV introduces an analytical model for tuning power method. Section V gives results and Section VI concludes the paper.

II. RELATED WORK

The issue related to thermal variation in on-chip optical communication has been addressed at both device and system levels. At device level, solutions relying on athermal devices [6] [7], voltage tuning [8], local heating [9][10], thermal-aware MR synthesis [11], and feed-back

control schemes [12] have been explored to limit the thermal impact on or control the resonant wavelength of MRs. At system level, the analyses allow the influence of temperature variation on the optical signal power received at receiver side [13] [14] to be evaluated. In our previous work [16], we propose a thermal-aware methodology to design an Optical Network-On-Chip (ONoC). Dynamic Voltage Frequency Scaling (DVFS) and workload migration techniques can be applied to reduce the temperature variation [17]. Moreover, a thermally-aware job allocation policy [18] is used to minimize the temperature gradients among MRs. This work relies on device characteristics such as the Free Spectral Range (FSR), the number of waveguides and wavelengths. In our work, we propose a thermal-aware laser tuning method. We consider the influence of thermal variation on the BER and tuning power consumption, taking into account the actual efficiency of on-chip laser and chip activity. In [19], the authors explore the placement of shared on-chip lasers, which are located on a layer on top of the optical interconnect. In our contribution, we consider distributed CMOS-compatible on-chip lasers that are placed on the same layer with on-chip interconnect.

Channel hopping, along with additional MRs on the ends of the spectral range, is utilized for tuning in [20][22]. In [14], the thermal tuning is adopted to adjust the MR resonant wavelength to the neighboring channel wavelength. In order to counter-balance the temperature variation, communication channels can be remapped through ONoC reconfiguration [23].

In [24], the laser output power is reduced until BER requirements can still be reached for optical interconnects using on-chip lasers. The MRs wavelength mismatch is first compensated with electrical tuning; then (i.e. once the wavelengths are aligned), the laser output power is reduced. However, the impact of the driver current onto the laser temperature is not taken into account, which will significantly impact i) the laser output power and ii) the emitted signal wavelength. Moreover, the optical crosstalk for the signal is not taken into consideration, which may significantly impact the BER.

In our work, we propose to jointly tune MRs and lasers wavelengths, which results in a global reduction of the optical channel power consumption (i.e. including laser and MRs), taking into account the temperature of the optical devices and the crosstalk, along with the laser output power change. Our method is also complementary to channel remapping, which allows the remapping of MRs wavelength to the closest (from an energy point of view) laser wavelength.

III. THERMAL SENSITIVITY OF NANOPHOTONIC INTERCONNECTS

This section presents communication channels in Wavelength Division Multiplexing (WDM)-based optical interconnect and illustrates the problem we tackle in this paper. Finally, our proposed method is illustrated.

A. Optical Communications based on WDM

WDM-based silicon photonic interconnects allow multiple wavelengths to be utilized for communication channels. Figure 1 illustrates a MWSR-like architecture with on-chip lasers such as PCM-VCSELs, utilizing N_{WL} wavelengths ($\lambda_0, \lambda_1, \dots, \lambda_{N_{WL}-1}$). The light vertically emitted by the laser is coupled into a waveguide by using a taper (not shown in the figure for the sake of clarity). Multiple laser wavelengths are combined together into a single waveguide with a multiplexer (“MUX” for short in Figure 1), which can be implemented by a multimode interference (MMI) coupler [21]. The wavelengths are equally distributed in order to reduce the crosstalk.

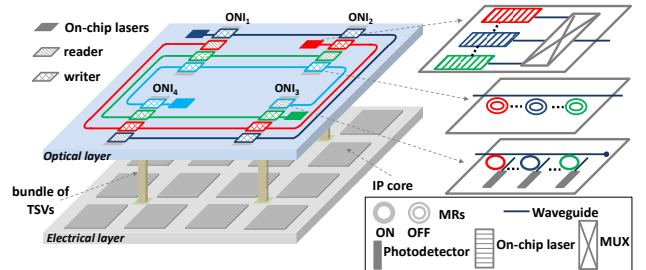


Figure 1: MWSR based optical interconnect with on-chip lasers.

B. Network Interface and Thermal Variation

The Optical Network Interfaces (ONIs) are responsible for modulating and receiving the optical signal on the optical layer, as shown in Figure 1. The optical signal (supplied by on-chip lasers) is modulated by one of the writers and propagates to the reader for detection and conversion. An interface is composed of i) writers for modulating signal and ii) one reader for detecting and receiving signal. The writer is composed of N_{WL} MRs. The reader uses passive MRs and photodetectors to filter and detect the expected signal respectively.

In the architecture example illustrated in Figure 1, we assume 4 interfaces, meaning that there are 4 MWSR channels. In this example, we also assume a single waveguide per channel. There are thus 3 writers and 1 reader per channel, each writer being composed of N_{WL} modulators and the reader being composed of N_{WL} photodetector/passive MR couples. For proper communications, it is mandatory that the wavelengths of the laser sources are well aligned with the resonant wavelengths of the MRs (in the writers) and those of the passive MR (in the reader). However, silicon based optical devices are sensitive to thermal variation ($0.1\text{nm}/^\circ\text{C}$ typically [17]), which induces the mismatch of the wavelengths of silicon photonic devices along one communication channel. As a result, the optical signal power received at the reader side decreases, which results in a lower SNR and a higher BER. In addition, the efficiency of on-chip lasers is reduced when the temperature increases [16], i.e., the BER is further degraded. In order to deal with this issue, we propose a thermal-aware tuning method to ensure a targeted BER at the reader side while minimizing the power consumption.

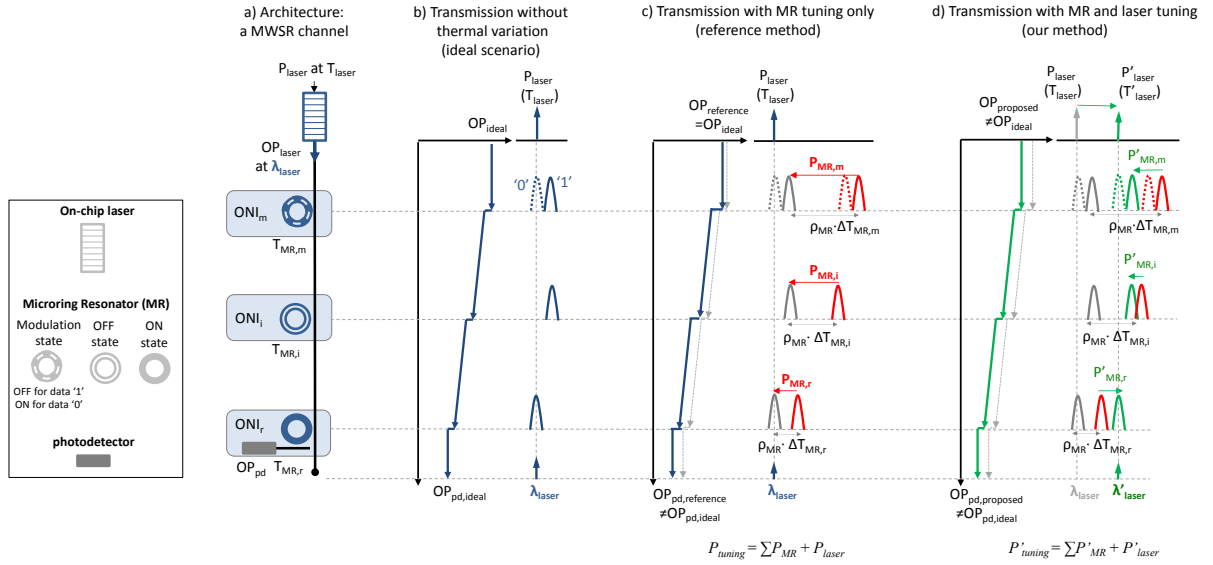


Figure 2: a) MWSR channel with one wavelength and two writers, b) transmission without thermal variation (ideal case), c) transmission considering MR tuning only method with off-chip laser sources, and d) transmission with our proposed on-chip laser tuning method.

C. Signals and MRs Wavelengths Alignment

Figure 2 illustrates our method in the context of a MWSR channel [2] with a single laser source, 2 writers and one reader. In our work, we consider distributed lasers that are located in the same layer with MRs, waveguide and photodetectors. For this purpose, we assume the use of CMOS-compatible laser sources such as VCSELs, for which the size is similar to the one of MRs [16]. As illustrated in Figure 2-a, ONI_m communicates with ONI_r : the MR in the intermediate ONI_i is turned OFF while the MR in ONI_m is in the modulation state (state OFF and ON to modulate data '1' and '0' respectively).

Figure 2-b illustrates the ideal transmission of a data '1' that occurs when there is no temperature variation along the communication path. The optical signal injected by the laser (which is characterized by a power OP_{ideal} and a wavelength λ_{laser}) crosses ONI_m and ONI_i , propagates along the waveguide until ONI_r where it is dropped to the photodetector (as illustrated by the blue transmission lines in the Figure 2-b). The power of the optical signal decreases along the path due to the waveguide propagation losses and the MRs crossing losses (blue line in Figure 2-b). The BER is estimated based on the received optical power $OP_{pd,ideal}$, the receiver sensitivity and the crosstalk (induced by other transmitting signals at different wavelengths, not illustrated here for the sake of clarity but taken into account in our models).

1) Existing MR Tuning Only Method

In case of temperature gradient over the communication path, the resonant wavelength of the MRs will drift (see the red transmission lines in Figure 2-c) while the wavelength of the emitted signal (λ_{laser}) remains the same in case off-chip lasers are considered. Without compensating the effect of this drift, the misalignment between the signal wavelength and the MRs resonant wavelengths leads to significantly increased BER. To overcome this effect, the MRs along the waveguide are tuned back to

their initial positions (the grey transmission lines in Figure 2-c) using thermal tuning or voltage tuning. The post-tuning signal transmission is illustrated by the blue line in Figure 2-c: the received optical power is slightly lower than in the ideal scenario due to marginal wavelengths misalignment. The MRs tuning power depends on their temperature drift ($\Delta T_{MR,m}$, $\Delta T_{MR,i}$ and $\Delta T_{MR,r}$) and the thermal sensitivity coefficient ρ_{MR} . The total power consumption of the channel is given by the sum of the laser power consumption P_{laser} and the MRs tuning power $\sum P_{MR}$.

2) Proposed Laser and MR Tuning Method

The key novelty of our method relies on the laser bias current tuning: since the laser temperature varies with its bias current, the wavelength of the optical signal can be tuned. Hence, in addition to tuning the MRs to align the resonant wavelengths with the optical signal, we also tune the wavelength of the optical signal itself, which contributes to reduce the power required to compensate the thermal variation.

As illustrated in Figure 2-d, tuning the driver current I_{laser} has an impact on the laser power consumption and the emitted signal wavelength: the former varies from P_{laser} to P'_{laser} and the latter shifts from λ_{laser} to λ'_{laser} . Therefore, under the same temperature gradient considered in Figure 2-c, the MRs wavelengths need to be tuned to λ'_{laser} , instead of λ_{laser} (see the green transmission lines in Figure 2-d). The MRs tuning power decreases since the wavelengths distance is reduced. As a drawback, the power of the emitted signal is reduced, meaning that a tradeoff needs to be defined in order to reach the target BER while decreasing the total power consumption of the channel. It is worth noticing that, although the method is illustrated with a 1-wavelength channel, it is generic and it can be applied to WDM channels, as illustrated in the results section. Moreover, it is complementary to related methods providing channel remapping [14][15] that we have adapted to minimize the tuning power instead of the

tuning distance.

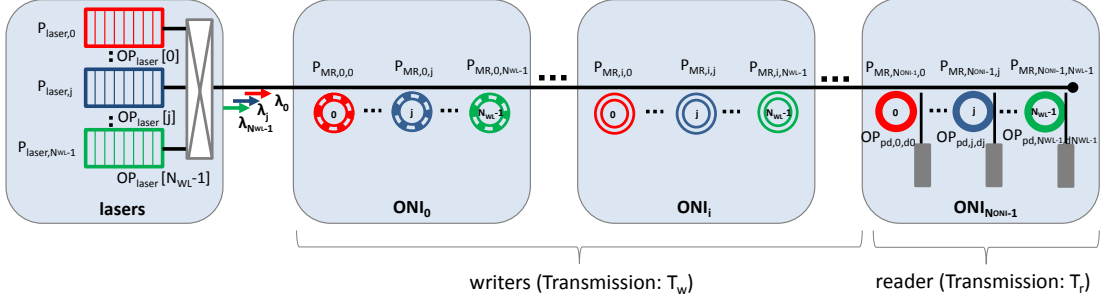


Figure 3: Generic MWSR channel.

IV. ANALYTICAL MODELS

This section presents the analytical models used for BER and tuning power estimations.

A. MWSR Channel

Figure 3 illustrates a generic MWSR channel with lasers and N_{ONI} interfaces ($N_{\text{ONI}}-1$ writers and 1 reader). N_{WL} lasers inject the optical signals, i.e., N_{WL} wavelengths can be used for communicating within the MWSR channel. In the figure, the signal modulation occurs in ONI_0 (i.e., the first writer). In the following, we define analytical models to evaluate i) the MWSR channel communication quality through a BER estimation and ii) the tuning power consumption (P_{tuning}) required to align the wavelengths of the optical signals with the MRs.

B. Bit Error Ratio (BER) Models

In Figure 3, the optical signals modulated with data d_j propagate through the intermediate writers (i.e., ONI_i , $i=1, \dots, N_{\text{ONI}}-2$) until reaching the reader, i.e., the targeted ($N_{\text{ONI}}-1$)th ONI. The optical power received by the photodetector at λ_j (OP_{pd,j,d_j}) is composed of i) the expected optical signal at λ_j (OP_{signal,j,d_j}) and ii) the crosstalk ($OP_{crosstalk,k,d_j}$) from the other signals at λ_k (where $k=0, 1, \dots, N_{\text{WL}}-1$, and $k \neq j$). Following [28], the signal transmission is thus composed by i) the transmission through the writers (T_w) and ii) the transmission through the reader part (T_r). Hence, the received optical power is defined as:

$$OP_{pd,j,d_j} = \sum_{k=0}^{N_{\text{WL}}-1} (T_{w,d_k} [k] \cdot T_{r,j} [k] \cdot OP_{laser} [k]) \quad (1)$$

The expected received signal power is:

$$OP_{signal,j,d_j} = T_{w,d_j} [j] \cdot T_{r,j} [j] \cdot OP_{laser} [j] \quad (2)$$

And the received crosstalk is:

$$OP_{crosstalk,k,d_j} = \sum_{k=0}^{N_{\text{WL}}-1} (T_{w,d_k} [k] \cdot T_{r,j} [k] \cdot OP_{laser} [k]) \quad (k \neq j) \quad (3)$$

We consider the signal at wavelength λ_k received by the j th photodetector as a general case. The transmission in the writers part ($T_{w,d_k} [k]$) and the reader part ($T_{r,j} [k]$) are detailed as following equations.

$$T_{w,d_k} [k] = (L_b)^{N_{w,b}[k]} (L_p)^{I_{r,total}[k]} \left(\prod_{n=0}^{N_{\text{ONI}}-1} \varphi_t(\lambda_k, \lambda_n + \Delta\lambda) \right)^{N_{w,r}-1} \left(\prod_{n=0}^{N_{\text{WL}}-1} \varphi_t(\lambda_k, \lambda_n + d_n \cdot \Delta\lambda) \right) \quad (4)$$

$$T_{r,j} [k] = (L_b)^{N_{r,b}[k]} (L_p)^{I_{r,total}[k]} \prod_{n=0}^{j-1} \varphi_t(\lambda_k, \lambda_n) \varphi_d(\lambda_k, \lambda_j) \quad (5)$$

Here, φ_t and φ_d are the signal transmission on MR

through port and drop port [28]. L_b and L_p are the bending loss and the propagation loss, which are assumed to be 0dB and 0.5dB/cm in this work. $I_{w,total}[k]$ and $I_{r,total}[k]$ (resp. $N_{w,b}[k]$ and $N_{r,b}[k]$) are the total waveguide length (resp. number of waveguide bends) experienced by signal at λ_k in writers and reader parts. $\Delta\lambda$ is the MR wavelength shift between ON and OFF states.

On the reader side, the receiver sensitivity gives the minimum optical signal power that can be detected by a photodetector. For a given receiver sensitivity, the SNR can be calculated as follow.

$$SNR = \frac{\mathfrak{R} \cdot (OP_{pd,j,1} - OP_{pd,j,0})}{i_n} \quad (6)$$

where i_n is the photodetector internal noise (4uA [25] in this work), \mathfrak{R} is the responsivity of the photodetector (1A/W in this work). $OP_{pd,j,1}$ and $OP_{pd,j,0}$ is the optical power received for the modulation of data "1" and "0" respectively. The worst-case SNR is estimated by considering

$$OP_{pd,j,1} = OP_{signal,j,1} \quad (7)$$

$$OP_{pd,j,0} = OP_{crosstalk,k,0} \quad (8)$$

The BER is obtained as following:

$$BER = \frac{1}{2} \operatorname{erfc} \left(\frac{SNR}{2\sqrt{2}} \right) \quad (9)$$

C. Power Models

The total tuning power consumption of a MWSR channel is divided into two parts: i) the lasers power consumption (P_{laser}) and ii) the MRs tuning power (P_{MR}) as in equations (10), (11) and (12). Table 1 summarizes the parameters involved in the model and the following details the lasers and MR models we assume.

$$P_{\text{tuning}} = P_{MR} + P_{laser} \quad (10)$$

$$P_{MR} = \sum_{i=0}^{N_{\text{ONI}}-1} \sum_{j=0}^{N_{\text{WL}}-1} P_{MR,i,j} \quad (11)$$

$$P_{laser} = \sum_{j=0}^{N_{\text{WL}}-1} P_{laser,j} \quad (12)$$

1) PCM-VCSEL for On-Chip Laser Sources

In this work, we consider the PCM-VCSELs (illustrated in Figure 4-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 [26]. 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 [26].

TABLE 1: TECHNOLOGICAL PARAMETERS

Parameters	Description	Unit
N_{ONI}	Number of ONIs in the network	
N_{VL}	Number of wavelengths per interface	
$P_{MR,i,j}$	Tuning power of the MR _i in ONI _i	mW
$P_{laser,j}$	Power consumption of laser _j	mW
$TE_{i,j}$	Tuning efficiency (voltage tuning or thermal tuning) of the MR _i in ONI _i	mW/n
$\lambda_{laser,j}$	Wavelength of laser corresponding to the wavelength _j after tuning	nm
$\lambda_{MR,i,j}$	Wavelength of MR _j in ONI _i considering the wavelength drift	nm
$\Delta\lambda_{laser,j}$	Wavelength drift of laser _j	nm
$\lambda_{laser,j,room}$	Wavelength of laser _j at room temperature	nm
$\lambda_{MR,i,j,room}$	Wavelength of MR _j in ONI _i at room temperature	nm
ρ_{MR}, ρ_{laser}	Temperature sensitivity factor of MR and laser wavelengths	nm/°C
$\Delta T_{MR,i,j}$	Temperature drift of MR _j in ONI _i and laser _j	°C
$\Delta T_{laser,j}$	Temperature drift of laser _j	°C
$T_{MR,i,j}$	Temperature of the MR _j in ONI _i and laser _j	°C
$T_{laser,j}$	Temperature of laser _j	°C
T_{room}	Room temperature as reference	°C
FSR	Free Spectral Range	nm
m	Resonant mode number of MR	

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. We assume an 80% coupling efficiency, which is slightly pessimistic compared to the 85% simulated in [27]. The signal wavelength (λ_{laser}) can be tuned by changing the laser temperature (T_{laser}) since the laser is sensitive to the thermal variation (assumed as 0.1nm/°C [26]), as shown in Figure 4-b. Under a given driver current (I_{laser}), the laser efficiency will decrease with an increase of the temperature (T_{laser}), which leads to a reduction of the emitted optical signal (OP_{laser}), as shown in Figure 4-c. The relationship between the wavelength and the temperature is given as:

$$\lambda_{laser,j} = \lambda_{laser,j,room} + \Delta\lambda_{laser,j} \quad (13)$$

$$\Delta\lambda_{laser,j} = \rho_{laser} \cdot \Delta T_{laser,j} \quad (14)$$

$$\Delta T_{laser,j} = T_{laser,j} - T_{room} \quad (15)$$

Since the lasers are located above the processing layer, their temperature is influenced by i) the driver current (I_{laser}) and ii) the chip activity. The laser power consumption ($P_{laser,j}$ for a laser at λ_j) can be estimated as:

$$P_{laser,j} = [slop_{ohm} \times (I_{laser,j} - I_{th}) + 1] \times I_{laser,j} - slop_{W/A} \times (I_{laser,j} - I_{th}) \quad (16)$$

where $slop_{ohm}$, $slop_{W/A}$ and I_{th} are the laser voltage slope, the output power slope and the threshold current, respec-

tively.

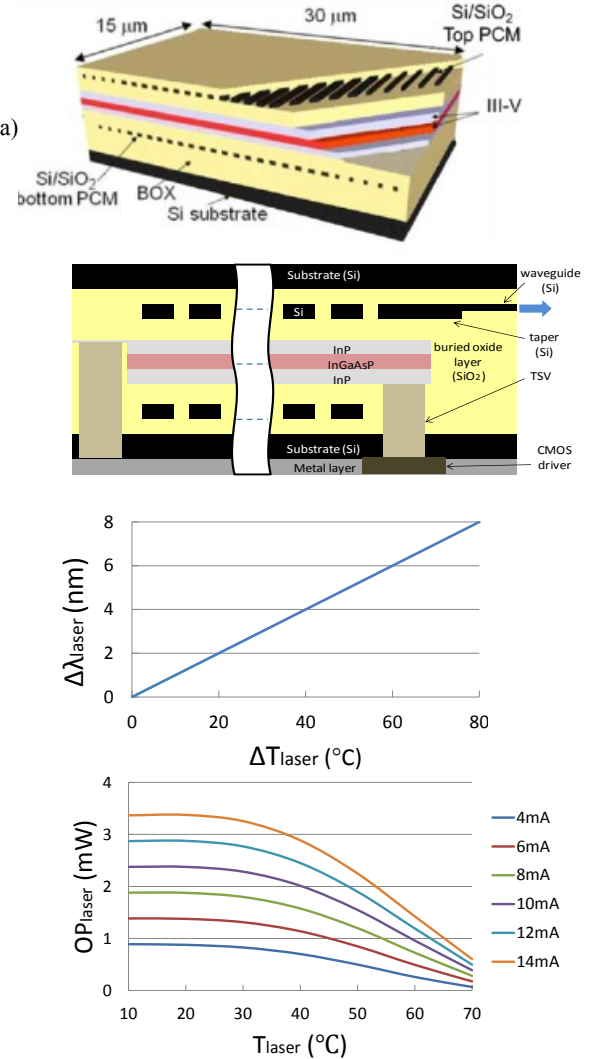


Figure 4: PCM-VCSEL: a) 3D view extracted from [26] and the cross-section view including the taper, b) wavelength drift of VCSEL ($\Delta\lambda_{laser}$) according to temperature drift (ΔT_{laser}) considering temperature sensitivity (ρ_{laser}) as 0.1nm/°C, and c) output power of laser (OP_{laser}) wrt. T_{laser} and I_{laser} . Curves b) and c) are obtained by extrapolation of data from [26].

2) Microring Resonators (MRs)

The modulation is realized by electro-optic effect on the MRs. Forward biased is applied to perform voltage tuning, which leads to a blue shift of the resonant wavelength. MRs can be classified according to their junction, PN or PIN. With a PN junction, the reverse biasing changes the refractive index through carrier depletion while, with a PIN junction, the refractive index is changed by carrier injection. Although PN junctions allow a faster switching time compared to PIN junctions, the higher extinction ratio provided by PIN junction leads to better communications [28]. In this work, we thus consider the use of PIN junctions only.




Figure 5: Top-view schematic of a PIN MR-based modulator with integrated voltage tuner and thermal tuner, inspired from [31].

Since the MRs are sensitive to temperature, their resonant wavelength needs to be tuned to ensure a proper alignment of the signal wavelength $\lambda_{laser,j}$ with the MR wavelength $\lambda_{MR,i,j}$. The MR tuning power ($P_{MR,i,j}$) depends on the distance between the wavelengths. MR tuning can be achieved by using electro-optic and thermo-optic effects [31], as illustrated in Figure 5 (monitor and feedback-control parts are not shown). The following details the tuning methods:

$$P_{MR,i,j} = TE_{i,j} \cdot |\lambda_{laser,j} - \lambda_{MR,i,j}| \quad (17)$$

$$\lambda_{MR,i,j} = \lambda_{MR,i,j,room} + \rho_{MR} \cdot \Delta T_{MR,i,j} \quad (18)$$

$$\Delta T_{MR,i,j} = T_{MR,i,j} - T_{room} \quad (19)$$

Voltage Tuning (VT) for blue shift: Voltage tuning is fast but its range (VTr) is limited to 1nm [17]. We denote VT_e as the voltage tuning efficiency (typically 0.13mW/nm [22][35]).

Thermal Tuning (TT) for red shift: the MR resonant wavelength can be red shifted by using local micro-heater [29]. Thermal tuning is slower than voltage tuning but its operating range (TTr) can reach 20nm [30]. The thermal tuning efficiency TT_e is lower (typically 0.24mW/nm [22][35]).

D. Simulation Model

From the estimation of the tuning power and the BER, we define exploration strategies to i) reduce the power consumption by maintaining a targeted BER and ii) decrease the BER as much as possible and estimate the power cost. Both methods require tuning the laser bias current and will be further detailed in Section V-B.

To evaluate these methods, the temperature of the optical devices (i.e., on-chip lasers and MRs) along the communication channel and under a given chip activity is needed. For this purpose, *IcTherm*¹ [32] is used to provide thermal maps of the silicon photonic interconnect. *IcTherm* is a thermal simulator for electronic devices which accurately models their complex structure and provides 3D full-chip temperature maps. *IcTherm* solves the physical equations that govern the temperature in the chip, using the Finite Volume Method [33], a numerical method for solving partial differential equations. It was validated against the commercial simulator COMSOL [34]: its maximal error was found to be less than 1% [32]. The structure of the system is discretized into small cubic cells that match the distribution of the materials and the heat sources.

V. RESULTS

This section describes the system we consider and then evaluates the efficiency of the proposed method.

A. Case Study

Figure 6 shows the global view of the targeted system, which contains the following components: steel back-plate, motherboard, socket, Intel's Single-Chip Cloud Computer (SCC chip) with silicon-photonic interconnects and on-chip laser sources, copper lid and heat sink.




Figure 6: Packaging of the SCC chip and the optical interconnect.

The considered 3D architecture is based on SCC (Figure 7-a) with a stacked optical layer [16]. Figure 7-b shows the abstract layout of SCC we use for the considered architecture layout. Following the method described in [16], the silicon photonic interconnect is placed on top of the SCC chip and thermal simulations are performed using *IcTherm* [32]. From the resulting thermal map, the proposed laser tuning method is applied in order to i) evaluate the BER using the transmission models from [28] and ii) evaluate the total channel power consumption P_{tuning} (including the laser power consumption P_{laser} and the MR tuning power P_{MR}).

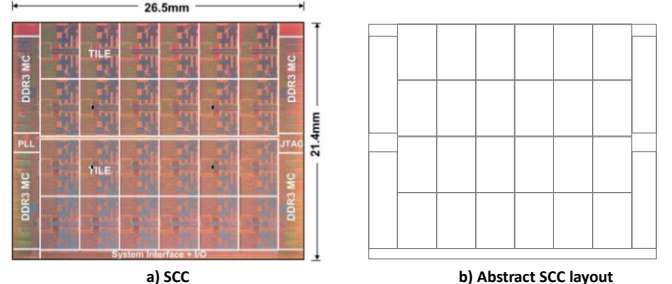


Figure 7: Considered electrical layer: a) SCC and b) its abstract layout.

¹ *IcTherm* website : <http://www.ictherm.com/>

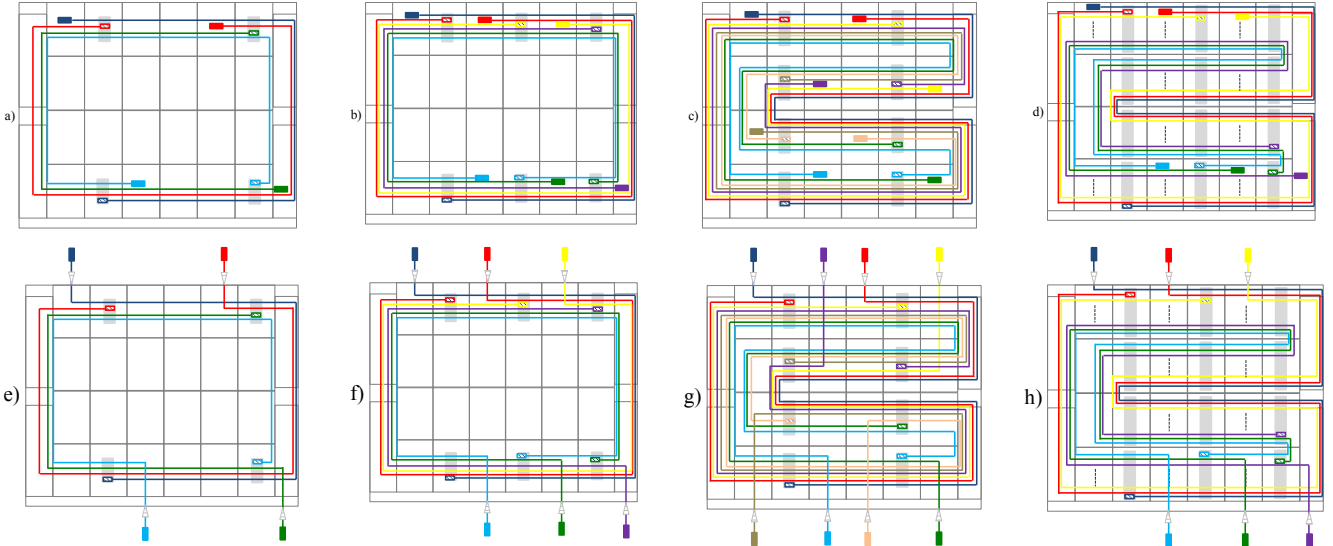


Figure 8: Considered interconnect layouts including 4, 6, 8 and 12 interfaces with a-d) on-chip lasers and e-h) off-chip lasers.

We assume 4 architectures including 4, 6, 8 and 12 interfaces, as illustrated in Figure 8. The number of interfaces gives the number of MWSR channels. In the figure, a single waveguide per channel is represented and only parts of the channels are illustrated for the 12 interfaces architecture for the sake of clarity. For comparison purpose, we assume architectures with on-chip lasers (Figure 8 a-d) and off-chip lasers (Figure 8 e-h). On-chip lasers are placed out of the interface in order to limit the temperature variation of the MRs as the laser driver current is tuned. The off-chip lasers are placed all around the chip in order to reduce the additional-insertion losses (i.e., from waveguide propagation and waveguide crossing).

Figure 9 illustrates photonic layer temperature maps for the layout given in Figure 8-a, assuming 4 waveguides per MWSR channel and 16 wavelengths per waveguide

(i.e., there is a total of 64 laser sources per channel). DDR3 memory controllers and PLL have been shut down to highlight the impact of chip activity and laser power consumption on the temperature. In order to illustrate the variety of application to be executed, we consider three values for the laser power consumption (4mW, 8mW, and 12mW) and three chip activities (uniform 10%, diagonal 25%-50% and corner 50%-5%). This figure highlights the important thermal gradient in the lasers region. It is worth noticing that, by considering the VCSELs described in Figure 4, only very limited lasing effect is obtained for temperature above 80°C. Such scenario occurs for a 50% local activity and 12mW laser power consumption (see the red hotspot region).

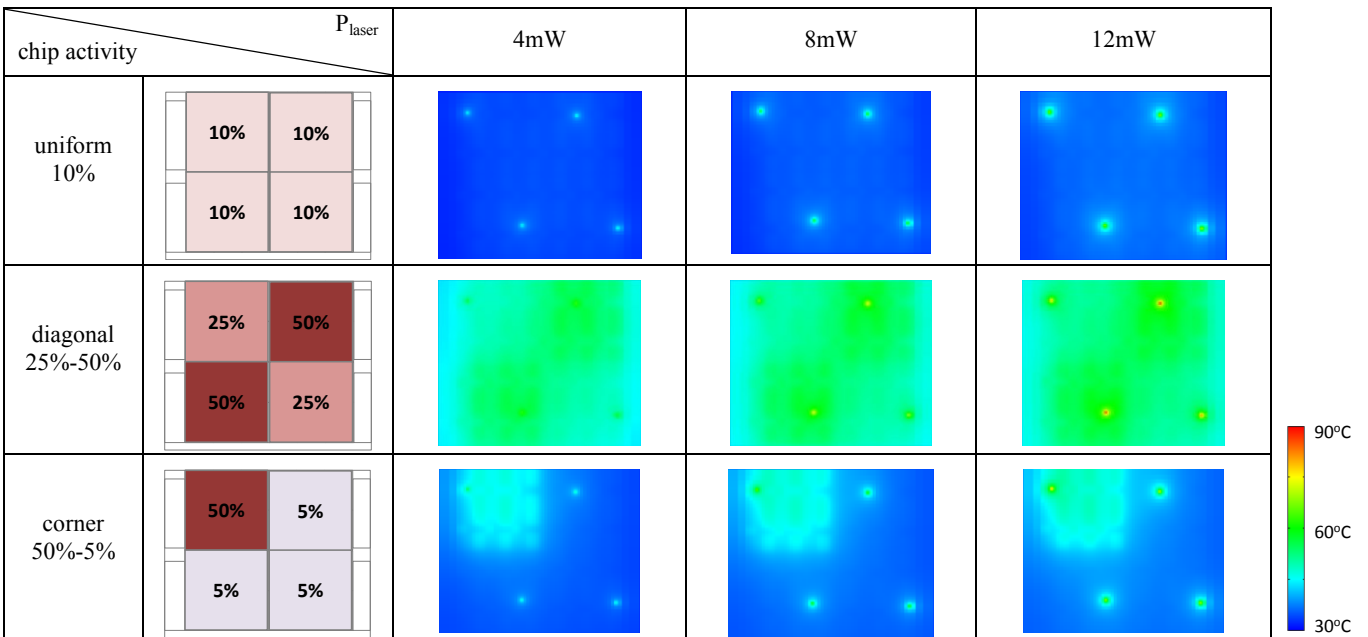


Figure 9: 2D temperature maps of the photonic layer for various chip activities (i.e., uniform 10%, diagonal 25%-50% and corner 50%-5%) and laser power consumption (i.e., 4mW, 8mW, and 12mW).

B. Impact of the Laser Bias Current on the BER and the Power Consumption

We evaluate our method for a 12 interfaces architecture (i.e., 12 MWSR channels with 11 writers and 1 reader per channel), 1 waveguide and 4 wavelengths (i.e., 4 lasers per waveguide and 4 MRs per writer). The study focuses on a single MWSR channel. We run thermal simulations under uniform 15% and 20% chip activity (i.e., 18.75W and 25W, respectively), with the tuning current I_{laser} ranging from 0mA to 16mA. Figure 10 represents the laser temperature drift ΔT_{laser} and the optical output power OP_{laser} deduced from the resulting thermal maps. It is worth noticing that, compared to Figure 4-c, a much higher sensitivity to the driver current is obtained due to the locally dissipated energy, which contributes to the temperature elevation. Indeed, Figure 4-c corresponds to characterization results obtained under a temperature stabilized using a Peltier cooling system. In Figure 10, the laser response is simulated in a 3D integrated circuit without cooling system.

For $I_{\text{laser}}=0.05\text{mA}$ under uniform 15% chip activity, $\Delta T_{\text{laser}} = 26^\circ\text{C}$ and there is no lasing effect (i.e. $OP_{\text{laser}}=0\text{mW}$) since the lasing current threshold for the corresponding temperature is not reached. Lasing effect starts at $I_{\text{laser}}=1\text{mA}$ and a maximum of 0.65mW optical power is obtained for 8mA. For this current value, the laser drift temperature is 36.8°C . Above 8mA, the laser efficiency drastically decreases due to the temperature elevation. From the laser only point of view, the most relevant I_{laser} current values ranges from 1mA to 8mA. However, from the whole MWSR channel point of view, current values above 8mA can also be relevant since a drift of the signal wavelengths accompanies the laser temperature elevation. Depending on the context, this may allow the reduction of the power needed to tune the MRs on the MWSR channel. For 20% chip activity, a similar trend is observed but ΔT_{laser} is slightly higher, which in turns leads to a reduction of the outputted light.




Figure 10: Impact of I_{laser} on the optical output power (OP_{laser}) and the laser temperature drift (ΔT_{laser}) under uniform 15% and 20% chip activities.

1) Results for Uniform Chip Activities

From the thermal maps over the optical layer, with a laser power consumption P_{laser} ranging from 1mW to 22mW (obtained by tuning the laser injection current I_{laser}), we evaluate the total power required to align the MRs resonant wavelengths with the emitted signals. We also evaluate the worst-case BER among all the channels in the network, as illustrated in Figure 11. For 20% chip activity

(in red), for small laser injection current (i.e., P_{laser} ranging from 1mW to 6mW), the optical power of the emitted signal is too low to compensate the channel losses, which results in a high BER. Then, the communication quality improves as the laser injection current increases: the BER reaches its optimal value at $P_{\text{laser}}=9\text{mW}$. By considering a static design method, the corresponding injection current will be selected to ensure the best communication quality. Above 9mW, the laser efficiency decreases due to the increase of the local temperature, leading to a higher BER. The channel power consumption is also given in the figure. It is composed of the laser power consumption (simply reported from the x-axis) and the MRs tuning power.

Figure 11 gives the tuning power consumption and the BER under 20% chip activity (in red). When the chip activity drops to 15% (in blue), two methods can be followed:

- To achieve a given BER (e.g., 10^{-12}), the laser current can be reduced to minimize the tuning power, as represented by a \rightarrow b arrow in Figure 11. In this example, the tuning power drops from 10.1mW to 7.5mW, i.e., 26% reduction.
- To minimize the BER, the laser current is slightly increased, as illustrated by c \rightarrow d arrow in Figure 11. In this example, the tuning power increases from 12.9mW to 14.4mW (i.e., 12% power cost). However, the BER remains optimal and decreases from 10^{-16} to 10^{-28} (for 20% and 15% chip activity respectively).

As reflected in the thermal maps, the gradient temperature within each interface remains below 1°C . The total tuning power of each MWSR channel can thus be obtained.

Simulations results show that the targeted 10^{-12} BER couldn't be reached for chip activity higher than 25%, independently from the selected laser power consumption. This limitation is directly related to the thermal sensitivity of the laser, for which the efficiency drops for temperature above 40°C . The laser efficiency is further explored in Section V-C.

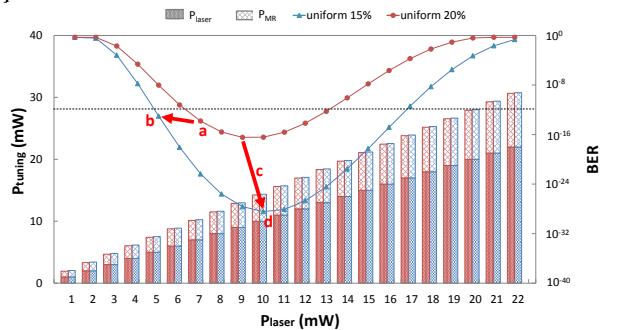


Figure 11: Tuning power (P_{tuning}) and BER for a MWSR channel with 12 interfaces, 4 wavelengths per waveguide under 15% and 20% uniform chip activities. FSR=59nm, $m=17$, voltage tuning efficiency $VTe=0.13\text{mW/nm}$, thermal tuning efficiency $TTe=0.24\text{mW/nm}$ [22][35].

2) Results for Diagonal Chip Activities

We also run thermal simulations for 5%-25% and 25%-5% diagonal chip activities. Figure 12 gives the results for

the MWSR channel with laser sources located on the top-left hand side of the chip (i.e., blue channel in Figure 8-d). 5mW and 6mW laser power consumption allow the total power consumption (while still achieving 10^{-12} BER) to be minimized for 5%-25% and 25%-5% activities, respectively. As expected, the maximum reachable BER is better for 5%-25% activity due to the lower heat locally dissipated by the chip (i.e., the optical signals outputted by the lasers is higher). An interesting trend is the slightly lower MRs tuning for 5%-25% activity, which is due to a reduced distance between the MRs resonant wavelengths and the laser wavelengths. While this power saving only slightly influences the total channel energy figures, significant gains are reachable for larger chip activity gradient, which is evaluated in the following.

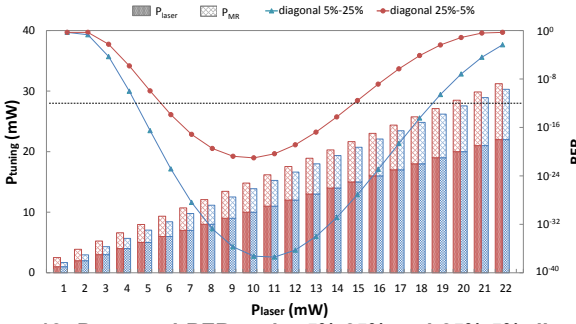


Figure 12: P_{tuning} and BER under 5%-25% and 25%-5% diagonal activities. Interconnect and technologies assumptions are those used in Figure 11.

3) Results for Corner Chip Activities

Figure 13 illustrates the total tuning power and the BER under a) 15%-65%, b) 15%-75% and c) 15%-85% corner activities. In this scenario, we only consider the MWSR channel with laser sources located in the 15% chip activity region since an insufficient lasing effect is obtained for lasers located in the region with higher activity.

In Figure 13-a, for $P_{\text{laser}}=1\text{mW}$, 30mW are required to align the MRs resonant wavelengths with the, significantly distant, optical signals. Increasing P_{laser} to 2mW leads to a reduction of the wavelengths distance, which in turn helps reducing P_{MR} to 20mW. P_{MR} continues shrinking until wavelengths are aligned, which is obtained for $P_{\text{laser}}=5\text{mW}$ (mark *a* in the figure). This value also corresponds to the minimum total tuning power allowing the targeted 10^{-12} BER to be reached. The optimal BER is obtained for $P_{\text{laser}}=11\text{mW}$ (mark *b*). In Figure 13-b, the targeted BER is obtained starting from $P_{\text{laser}}=6\text{mW}$. However, for this value, the wavelengths distance to compensate is still significant and leads to $P_{\text{MR}}=4\text{mW}$. By increasing P_{laser} to 7mW (mark *c*), 3mW can be saved on the MRs tuning, which leads to a more power efficient solution. This corresponds to the scenario sketched in Figure 2. Due to a slightly higher temperature, the optimal BER is obtained for $P_{\text{laser}}=10\text{mW}$ (mark *d*). The same trend can be observed in Figure 13-c: $P_{\text{laser}}=8\text{mW}$ allows the BER requirements to be reached but $P_{\text{laser}}=9\text{mW}$ is a globally more power efficient solution. Furthermore, $P_{\text{laser}}=9\text{mW}$ also leads to the optimal BER (mark *e*), thus demonstrating the potential for adaptive laser tuning methods to maximize the energy efficiency of nanophotonic interconnects.

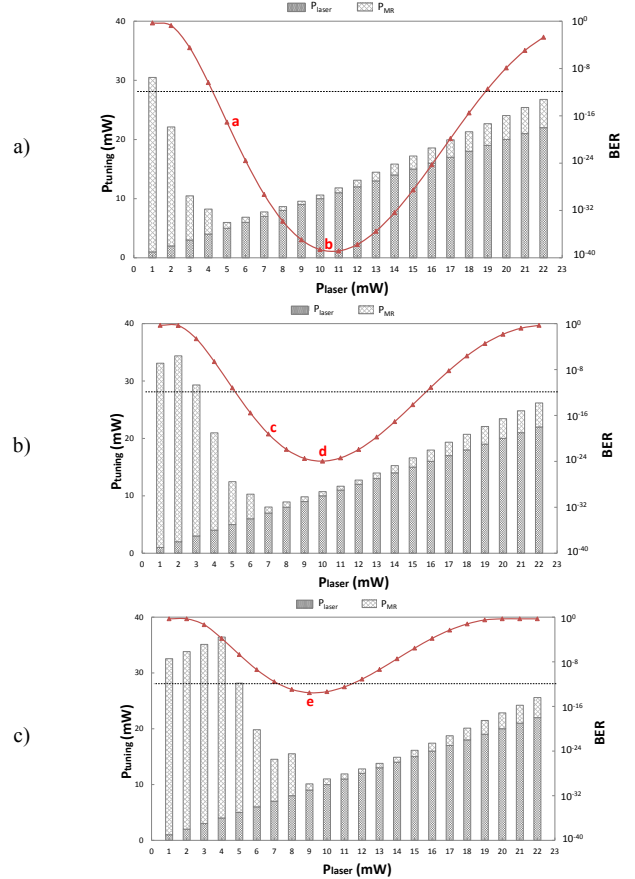


Figure 13: P_{tuning} and BER under a) 15%-65%, b) 15%-75% and c) 15%-85% corner chip activities. Interconnect and technologies assumptions are those used in Figure 11.

Finally, a slight 2% increase of the corner gradient activity (i.e., 15%-87%, not shown in the figure) will avoid reaching the targeted BER for $P_{\text{laser}}=11\text{mW}$ (which leads to the optimal BER in Figure 13-a, while $P_{\text{laser}}=9\text{mW}$ still achieves it. This demonstrates that for the considered system, an adaptive method relying on laser tuning not only reduces the energy, but also helps covering a wider range of chip activities.

C. Laser Efficiency Comparison

In this section, we consider two laser efficiencies and we evaluate their impact on the tuning power. Figure 14-a represents a conservative and an aggressive scenario. The conservative scenario corresponds to the laser already detailed in Figure 4 and the aggressive one is obtained by considering a x2 laser efficiency. The tuning power is illustrated in Figure 14-b for uniform chip activities ranging from 5% to 20%, where we assume 12 interfaces. The targeted BER is set to 10^{-12} , which requires less energy with aggressive scenario since the receiver optical power is higher for a same bias current.

Hence, for a 5% chip activity, P_{tuning} is reduced from 6.4mW (conservative) to 3.6mW (aggressive). The improvement of the laser efficiency leads to an increase of the ratio of the MR tuning in the channel total power consumption. For instance, for 5% chip activity, the MR tuning power ratio increases from 37% to 45%. However, due to the higher efficiency of the laser, tuning the laser

becomes more efficient than tuning the MRs, which leads to a reduction of the MRs tuning power (i.e., 2.36mW and 1.62mW for conservative and aggressive scenario respectively). The combined reduction of laser tuning and MRs tuning leads to significant global energy saving, e.g., 43% for 5% chip activity. Further power saving are obtained for higher chip activities (e.g., 54% under 20% chip activity). These results would help investigating laser requirements for ultra-low power silicon photonics interconnects.

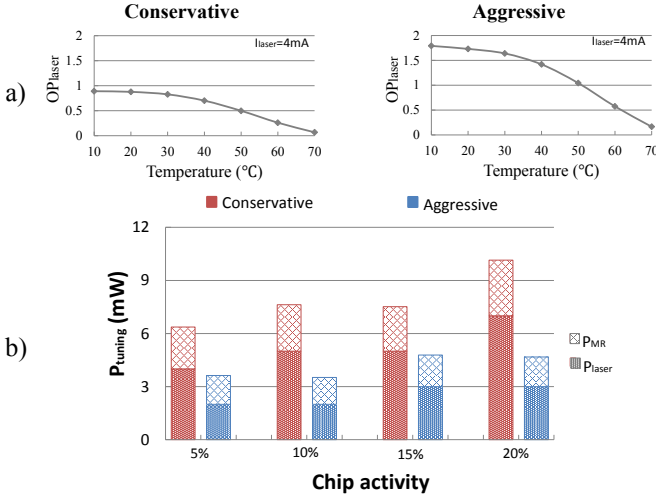


Figure 14: a) Laser efficiency for conservative and aggressive scenarios and b) P_{tuning} for 5%, 10%, 15% and 20% chip activities. Interconnect and technologies assumptions are those used in Figure 11.

D. Tuning Methods Comparison

We compare our method with related solutions in which only MR tuning is used. We consider related solutions relying on off-chip and on-chip lasers. For off-chip lasers, the laser wavelength is fixed and the resonant wavelengths of the MRs along a given channel may need to be tuned back to be aligned with the laser, considering channel remapping. We also take into account waveguide crossing losses (considered as 0.05dB per crossing), assuming a single waveguide is used per channel. We assume the same characteristics for off-chip and on-chip lasers and a 80% coupling efficiency has been considered, which is slightly optimistic compared to the 74% demonstrated in [36]. For on-chip lasers with MR tuning only method, we take into account the wavelength drift of the emitted optical signal due to the temperature variation to evaluate the energy needed to tune the MRs. With our method, the laser bias current I_{laser} is tuned.

We consider a 12 interfaces architecture with 4 to 16 wavelengths per MWSR channel and a 10Gb/s data rate. The initial chip activity is assumed to be 10% and the laser power consumption is set to achieve a given BER (e.g., 10^{-12}). Figure 15 illustrates the contributions of the laser and the MR tuning power (without considering the modulation) for 4 to 16 wavelengths. Our method is the most power efficient for all the studied cases. For instance, for 4 wavelengths MWSR channels, we reach 2.5pJ/bit wrt. 3.1pJ/bit and 10.7pJ/bit for related on-chip and off-chip methods respectively. When the number of wavelength increases, additional energy is needed since

the number of MRs considerably increases.

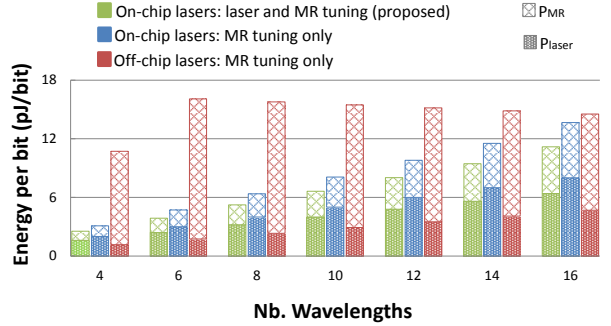


Figure 15: Channel energy efficiency w and w/o laser tuning for 12 interfaces. FSR=59nm, $m=17$, $VTe=0.13mW/nm$, $TTe=0.24mW/nm$ [22][35].

The results show that, for 16 wavelengths, the optical channels using off-chip lasers give comparable power results to our method. It is also worth noticing that the results obtained for the off-chip lasers method are optimistic for 3 reasons. First, we assume off-chip lasers working at a constant 25°C room temperature, which requires an energy consuming cooling system not taken into account here. Second, we assume a single waveguide per MWSR channel, which leads to at least 11 waveguide crossing for a 12 interfaces architecture illustrated in Figure 8-h (e.g., it leads to 0.55dB losses on each waveguide). MWSR channels with more than 4 waveguides are regularly considered in the literature, which would significantly increase the losses. Third, the propagation loss due to the longer waveguides has not been considered. However, additional scenarios, including application dependant chip activities, should be considered to further consolidate this comparison.




Figure 16: Design tradeoff between BER and energy efficiency for 12 interfaces.

To further compare our method with the MR tuning only method, we also evaluate for which number of wavelengths the BER is the best. For this purpose, we lower as much as possible the BER when the chip activity decreases from 20% to 15% and we evaluate the required energy. We assume 12 interfaces and a number of wavelengths ranging from 4 to 16. As shown in Figure 16, the energy/bit remains slightly constant with the number of wavelengths (approx. 1.2 pJ/bit). However, the BER significantly increases from 10^{-28} (4 wavelengths) to 10^{-10} (16 wavelengths), i.e., the communication quality is better for small number of wavelengths. Such method allows the maintaining of a given BER level to satisfy applications constrained by communication quality. Overall, these results show that the MWSR channels with a small number of wavelengths provide a better BER. This leads us to

conclude that our method tends to be the most power efficient for this range of wavelengths.

E. Energy Saving under a Given BER

We evaluate the potential gain of the proposed laser tuning method with the traditional on-chip laser (for which only MRs tuning is possible). For this purpose, we first set the laser power consumption to 9mW in order to reach a targeted 10^{-12} BER for the highest considered chip activity (25% in this example). We then reduce the chip activity to 20%, 15% and 10%, and we evaluate the achievable power reduction while still reaching 10^{-12} BER [37][38]. As illustrated in Table 2, up to 42% reduction is obtained for a 10% chip activity. In case a 10^{-9} BER [39][40] turns to be acceptable (i.e., in order to match the requirements of an application to be executed), the energy consumption can be further decreased. These results demonstrate the significant energy reductions achievable by using chip activity-aware methods.

TABLE 2: ENERGY CONSUMPTION REDUCTION: W/ LASER TUNING WRT. W/O LASER TUNING

Chip activity \ Target BER	20%	15%	10%
10^{-12}	21%	42%	42%
10^{-9}	32%	42%	52%

We evaluate the energy saving of our method for various numbers of interfaces and wavelengths and by assuming the same technique as previously described. As shown in Figure 17, the energy saving is the smallest for 12 interfaces and 16 wavelengths but it still reaches 31% wrt. the MRs tuning only method. On average, 53% energy saving is obtained.




Figure 17: Energy saving for our joint laser and MR tuning method wrt. MR tuning only method for different channel configurations.

VI. CONCLUSION

In this paper, we propose to jointly tune lasers and MRs to improve the energy efficiency of fully integrated nanophotonic interconnects. For this purpose, we have defined a method relying on thermal simulations and crosstalk analyses, taking into account the thermal sensitivity of the laser sources. Evaluations have been carried out for a 3D stacked architecture interconnecting processors with MWSR-like optical channels. We have assumed CMOS compatible PCM-VCSELs for the light sources and the layout of the optical layer has been designed to avoid

waveguide crossing. Design space exploration covered the number of interfaces, the number of wavelengths and the laser efficiency, taking into account uniform, diagonal and corner chip activities.

Compared to methods for which laser tuning is not possible, results show that a combined tuning of laser and MRs leads to 53% energy reduction when the uniform chip activity decreases from 20% to 5%. BER-energy tradeoffs have been explored and allow strategies to be defined to minimize either the energy, or the BER. As a key result, we have shown that, under specific chip activities, increasing the laser power consumption allows both energy and BER to be improved. This trend has been observed for a MWSR channel interconnecting 12 interfaces. We also showed that, by being able to tune the laser power within the 7-13mW ranges, the targeted 10^{-12} BER is reachable in all the scenarios we covered, assuming a maximum 20% chip activity in the region where the lasers are located. This strong limitation directly depends on the laser efficiency, which drastically decreases above 40°C. Major technological improvements are thus needed to make silicon photonics becoming a realistic and viable solution for on-chip interconnects in fully integrated 3D architectures. While such improvement would lead to more energy efficient optical channels by reducing the laser power consumption, it is worth noticing that our joint laser and MRs tuning method would further contribute to this objective by also reducing the MRs tuning power.

A major challenge for the practical use of our method at run-time is the calibration process, which has to rapidly explore lasers and MRs tuning options. For this purpose, transient thermal simulations are mandatory to optimize the controller, which will be investigated in the near future. In parallel, we are currently investigating whether layout optimization contributes to improve the energy efficiency of nanophotonic interconnects. Together with the run-time calibration process, one can define novel group tuning methods for lasers that could ideally complement the already existing group tuning for MRs. Finally, our method is generic and can be applied to other types of light sources, included on-chip and off-chip lasers. This leads to a need for comparisons we will achieve in the future.

ACKNOWLEDGMENTS

The authors would like to thank the editor and anonymous reviewers for their valuable feedback. The authors also gratefully acknowledge Xavier Letartre from INL lab for the fruitful discussions. Hui LI is supported by China Scholarship Council (CSC).


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] 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 *Proc. 35th Ann. Int. Symp. Computer Architecture (ISCA'08)*, 2008.
 - [3] G. Roelkens, L. Liu, D. Liang, R. Jones, A. Fang, B. Koch, and J. Bowers, "III-V/silicon photonics for on-chip and inter-chip optical interconnects," *Laser & Photon. Rev.*, Vol. 4, No. 6, pp. 751-779, 2010.
 - [4] C. Sciancalepore, B. B. Bakir, X. Letartre, J. Harduin, N. Olivier, C. Seassal, J.-M. Fedeli, and P. Viktorovitch, "CMOS compatible ultra-compact 1.55-um emitting VCSEL using double photonic crystal mirrors," *IEEE Photon. Technol. Lett.*, Vol. 24, No. 6, pp. 455–457, 2012.
 - [5] C. Sciancalepore, B. B. Bakir, X. Letartre, N. Olivier, D. Bordel, C. Seassal, P. Rojo-Romeo, P. Regreny, J.-M. Fedeli, and P. Viktorovitch, "CMOS-compatible integration of III-V VCSELs based on double photonic crystal reflectors," in *Proc. 8th IEEE Int. Conf. GFP*, pp. 205–207, 2011.
 - [6] S. S. Djordjevic, K. Shang, B. Guan, S. T. S. Cheung, L. Liao, J. Basak, H.-F. Liu, and S. J. B. Yoo, "CMOS-compatible, athermal silicon ring modulators clad with titanium dioxide," *OPTICS LETTERS*, Vol. 21, No. 12, 2013.
 - [7] M. Mohamed, Z. Li, X. Chen, L. Shang, and A. R. Mickelson, "Reliability-Aware Design Flow for Silicon Photonics On-Chip Interconnect," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, Vol. 22, No. 8, 2014.
 - [8] S. Manapatruni, R. K. Dokania, B. Schmidt, N. Sherwood-Droz, C. B. Poitras, A. B. Apsel, and M. Lipson, "Wide temperature range operation of micrometer-scale silicon electro-optic modulators," *OPTICS LETTERS*, Vol. 33, No. 19, pp. 2185-2187, 2008.
 - [9] A. Biberman, N. Sherwood-Droz, B. G. Lee, M. Lipson and K. Bergman, "Thermally Active 4x4 Non-Blocking Switch for Networks-on-Chip," in *21st Ann. Meeting IEEE Lasers and Electro-Optics Society (LEOS'08)*, 2008.
 - [10] C. Sun, C.-H. Chen, G. Kurian, L. Wei, J. Miller, A. Agarwal, L.-S. Peh, and V. Stojanovic, "DSSENT - A Tool Connecting Emerging Photonics with electronics for Opto-Electronic Networks-on-Chip Modeling," in *Sixth IEEE/ACM Int. Symp. Networks on Chip (NoCS'12)*, 2012.
 - [11] C. Condrat, P. Kalla, and S. Blair, "Thermal-aware Synthesis of Integrated Photonic Ring Resonators," In *Proc. IEEE/ACM Int. Conf. Computer-Aided Design (ICCAD'14)*, 2014.
 - [12] K. Padmaraju, J. Chan, L. Chen, M. Lipson, and K. Bergman, "Thermal stabilization of a microring modulator using feedback control," *Optics Express*, Vol. 20, No. 27, pp. 27999-28008, 2012.
 - [13] Y. Ye, J. Xu, X. Wu, W. Zhang, X. Wang, M. Nikdast, Z. Wang, and W. Liu, "System-Level Modeling and Analysis of Thermal Effects in Optical Networks-on-Chip," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, Vol. 21, No. 2, pp. 292-305, 2013.
 - [14] Y. Ye, Z. Wang, J. Xu, X. Wu, X. Wang, M. Nikdast, Z. Wang, and L. H. K. Duong, "System-Level Modeling and Analysis of Thermal Effects in WDM-Based Optical Networks-on-Chip," *IEEE Trans. Computer-Aided Design of Integrated Circuits and Systems*, Vol. 33, No.11, pp.1718-1731, 2014.
 - [15] M. Georgas, J. Leu, B. Moss, C. Sun, and V. Stojanovic, "Addressing Link-Level Design Tradeoffs for Integrated Photonic Interconnects," in *Proc. IEEE Custom Integrated Circuits Conference (CICC'11)*, 2011.
 - [16] H. Li, A. Fourmigue, S. Le Beux, X. Letartre, I. O'Connor, and G. Nicolescu, "Thermal Aware Design Method for VCSEL-based On-Chip Optical Interconnect," In *Proc. Design, Automation & Test in Europe Conference & Exhibition (DATE'15)*, 2015.
 - [17] Z. Li, M. Mohamed, X. Chen, E. Dudley, K. Meng, L. Shang, A. Mickelson, R. Joseph, M. Vachharajani, B. Schwartz, and Y. Sun, "Reliability Modeling and Management of Nanophotonic On-Chip Networks," *IEEE Trans. Very Large Scale Integration (VLSI) Systems*, Vol. 20, No. 1, pp. 98-111, 2012.
 - [18] T. Zhang, J. L. Abellán, A. Joshi, and A. K. Coskun, "Thermal management of Manycore Systems with Silicon-Photonic Networks," In *Proc. Design, Automation & Test in Europe Conference & Exhibition (DATE'14)*, 2014.
 - [19] C. Chen, T. Zhang, P. Contu, J. Klamkin, A. K. Coskun, A. Joshi, "Sharing and Placement of On-chip Laser Sources in Silicon-Photonic NoCs," in *Eighth IEEE/ACM Int. Symp. Networks-on-Chip (NoCS'14)*, 2014.
 - [20] M. Mohamed, Z. Li, X. Chen, L. Shang, A. Mickelson, M. Vachharajani, and Y. Sun, "Power-Efficient Variation-Aware Photonic On-Chip Network Management," in *ACM/IEEE Int. Symp. Low-Power Electronics and Design (ISLPED'10)*, 2010.
 - [21] F. Mandorlo, P. R. Romeo, N. Olivier, L. Ferrier, R. Orbotchouk, X. Letartre, J. M. Fedeli, and P. Viktorovitch, "Controlled Multi-Wavelength Emission in Full CMOS Compatible Micro-Lasers for on Chip Interconnections," *J. Lightwave Technol.*, Vol. 30, No. 19, pp. 3073-3080, 2012.
 - [22] C. Nitta, M. Farrens, and V. Akella, "Addressing System-Level Trimming Issues in On-Chip Nanophotonic Networks," In *Proc. IEEE 17th Int. Symp. High Performance Computer Architecture (HPCA '11)*, 2011.
 - [23] Y. Zhang, P. Lisherness, M. Gao, J. T. Bovington, K. T. Cheng, H. Wang, and S. Yang, "Power-Efficient Calibration and Reconfiguration for Optical Network-on-Chip," *J. Optical Communications and Networking*, Vol. 4, No. 12, pp. 955–966, 2012.
 - [24] R. Wu, C.-H. Chen, C. Li, T.-C. Huang, F. Lan, C. Zhang, Y. Pan, J. E. Bowers, R. G. Beausoleil, and K.-T. Cheng, "Variation-Aware Adaptive Tuning for Nanophotonic Interconnects," in *Proc. IEEE/ACM Int. Conf. Computer-Aided Design (ICCAD'15)*, pp 487-493, 2015.
 - [25] L. Vivien, A. Polzer, D. M.-Morini, J. Osmond, J. M. Hartmann, P. Crozat, E. Cassan, C. Kopp, H. Zimmermann, and J. M. Fédéli, "Zero-bias 40Gbit/s germanium waveguide photodetector on silicon," *Optics Express*, Vol. 20, No.2, pp. 1096–1101, 2012.
 - [26] C. Sciancalepore, B. Ben Bakir, C. Seassal, X. Letartre, J. Harduin, N. Olivier, J.-M. Fedeli, 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.
 - [27] 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.
 - [28] H. Li, S. Le Beux, Y. Thonnart, and I. O'Connor, "Complementary Communication Path for Energy Efficient On-Chip Optical Interconnects," In *Proc. 52nd Ann. Design Automation Conference (DAC '15)*, 2015.
 - [29] H. Shen, M. H. Khan, L. Fan, L. Zhao, Y. Xuan, J. Ouyang, L. T. Varghese, and M. Qi, "Eight-channel reconfigurable microring filters with tunable frequency, extinction ratio and bandwidth," *OPTICS EXPRESS*, Vol. 18, No. 17, 2010.
 - [30] F. Gan, T. Barwicz, M.A. Popović, M.S. Dahlem, C.W. Holzwarth, P.T. Rakich, H.I. Smith, E.P. Ippen and F.X. Kärtner,

"Maximizing the Thermo-Optic Tuning Range of Silicon Photonic Structures," in *IEEE/LEOS Photonics in Switching Conference*, 2007.

- [31] Y. Li and A. W. Poon, "Active resonance wavelength stabilization for silicon microring resonators with an in-resonator defect-state-absorption-based photodetector," *OPTICS EXPRESS*, Vol. 23, No. 1, 2015.
- [32] A. Fourmigue, G. Beltrame, and G. Nicolescu, "Efficient Transient Thermal Simulation of 3D ICs with Liquid-Cooling and Through Silicon Vias," In *Proc. Design, Automation & Test in Europe Conference & Exhibition (DATE'14)*, 2014.
- [33] S. C. Chapra and R. P. Canale. *Numerical Methods for Engineers*, McGraw-Hill, Inc., New York, NY, USA, 6th edition, 2009.
- [34] COMSOL. <http://www.comsol.com>, July 2014.
- [35] Y. Xu, J. Yang, and R. Melhem, "Tolerating Process Variations in Nanophotonic On-chip Networks," In *Proc. 39th Ann. Int. Symp. Computer Architecture (ISCA '12)*, 2012.
- [36] W. D. Sacher, Y. Huang, L. Ding, B. J. F. Taylor, H. Jayatilaka, G.-Q. Lo, and J. K. S. Poon, "Wide bandwidth and high coupling efficiency Si₃N₄-on-SOI dual-level grating coupler," *OPTICS LETTERS*, Vol. 22, No. 9, pp. 10938-10947, 2014.
- [37] M. Petracca, B.G. Lee, K. Bergman, and L. P. Carloni, "Design Exploration of Optical Interconnection Networks for Chip Multiprocessors," in *16th IEEE Symp. High Performance Interconnects (HOTI'08)*, 2008.
- [38] Z. Li, D. Fay, A. Mickelson, L. Shang, M. Vachharajani, D. Filipovic, W. Park, Y. Sun, "Spectrum: A Hybrid Nanophotonic-Electric On-Chip Network," in *46th ACM/IEEE Design Automation Conference (DAC '09)*, 2009.
- [39] Y. Xie, M. Nikdast, J. Xu, W. Zhang, Q. Li, X. Wu, Y. Ye, X. Wang, and W. Liu, "Crosstalk Noise and Bit Error Rate Analysis for Optical Network-on-Chip," In *Proc. 47th Design Automation Conference (DAC '10)*, 2010.
- [40] R. Ji, L. Yang, L. Zhang, Y. Tian, J. Ding, H. Chen, Y. Lu, P. Zhou, and W. Zhu, "Five-port optical router for photonic networks-on-chip," *OPTICS EXPRESS*, Vol. 19, No. 21, pp. 20258-20268, 2011.



Hui Li received the B.S. and M.S. in Telecommunications Engineering from Xidian University, China in 2010 and 2013, respectively. Since 2013, she has been pursuing her Ph.D. degree in the Heterogeneous System Design group at Ecole Centrale de Lyon, France. Her research interests include silicon photonics, optical interconnect, network on chip, reliability.



Alain Fourmigue received the M.Sc. and Ph.D. degrees in computer engineering from the Ecole Polytechnique of Montreal, Canada, in 2011 and 2014, respectively. He is the author of ICTherm, a fast thermal simulator for modern electronic devices. Currently, he is with Opal-RT Technologies Inc., Montreal, as a research engineer and software developer, specialized in real-time simulation for power systems. His current research interests include high-performance solvers for sparse linear systems, numerical methods on non-conformal meshes and model order reduction techniques.



Sébastien Le Beux is Associate Professor for Heterogeneous and Nanoelectronics Systems Design at Ecole Centrale de Lyon. He is currently responsible for nanoprocessors research activities at the Heterogeneous System Design group of the Lyon Institute of Nanotechnology CNRS laboratory. He obtained his PhD in Computer Science from the University of Sciences and Technology of Lille in 2007. His research interests include design methods for emerging (nano) technologies and embedded systems, including silicon photonic interconnect and reconfigurable architectures. He has authored or co-authored over 70 scientific publications including journal articles, book chapters, patent and conference papers. He serves on the steering committees, organizing committees, and technical program committees of international conferences such as DATE, CODES+ISSS, NOCS and NanoArch.



Ian O'Connor (IEEE S'95-M'98-SM'07) was born in Cambridge, U.K., in 1969. He received the European M.Sc. degree in electronics engineering from the University of Essex, Essex, U.K., in 1992, the Ph.D. degree in electronics from the University of Lille, Lille, France, in 1997, and the professoral dissertation (Habilitation à Diriger des Recherches) from École Centrale de Lyon, Ecully, France, in 2005. From 1993 to 1997, he was a Research and Teaching Assistant with the Institut Supérieur d'Electronique du Nord, Lille, France, and from 1997 to 1998, he was a Senior CAD Engineer with Philips Semiconductors, Southampton, U.K. He joined École Centrale de Lyon in 1998 as Associate Professor, where he is Professor for Heterogeneous and Nanoelectronics Systems Design and head of the Heterogeneous Systems Design group at the Lyon Institute of Nanotechnology. Since 2008, he also holds a position of Adjunct Professor at École Polytechnique de Montréal, Canada. His research interests include novel computing architectures based on emerging technologies, associated with methods for design exploration. He has authored or co-authored around 200 book chapters, journal publications, conference papers and patents and has presented invited papers at several major conferences.



Gabriela Nicolescu obtained her B. Sc. A and her MSc degree from Politehnica Bucharest. She obtained her PhD degree, in 2002, from INPG (Institut National Polytechnique de Grenoble) in France. She has been working at Ecole Polytechnique de Montréal (Canada) since 2003, where she is a professor in the Computer and Software Engineering Department. Her research interests are in the field of design methodologies, programming and security for systems with advanced technologies, such as 3D multi-processor systems-on-chip integrating liquid cooling and optical networks. She published five books and she is the author of more than hundred articles in journals, international conferences and book chapters.