CHAMELEON: CHANNEL Efficient Optical Network-on-Chip

Sébastien Le Beux, Hui Li, Ian O'Connor, Kazem Cheshmi, Xuchen Liu,
Jelena Trajkovic, Gabriela Nicolescu

► To cite this version:

Sébastien Le Beux, Hui Li, Ian O'Connor, Kazem Cheshmi, Xuchen Liu, et al.. CHAMELEON: CHANNEL Efficient Optical Network-on-Chip. IEEE International Conference on Design Automation and Test in Europe (DATE), Mar 2014, Dresden, Germany. 10.7873/DATE.2014.317 . hal-01075358

HAL Id: hal-01075358

<https://hal.inria.fr/hal-01075358>

Submitted on 17 Oct 2014

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.

CHAMELEON: CHANNEL Efficient Optical Network-on-Chip

Sébastien Le Beux^{1*}, Hui Li¹, Ian O'Connor¹,

Kazem Cheshmi², Xuchen Liu¹, Jelena Trajkovic², Gabriela Nicolescu³

¹ Lyon Institute of
Nanotechnology, INL-UMR5270
Ecole Centrale de Lyon,
Ecully, F-69134, France

² Electrical and Computer
Engineering Department
Concordia University
Montreal, QC, Canada

³ Computer and Software
Engineering Dept.
Ecole Polytechnique de Montréal
Montréal (QC), Canada

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

Abstract— The next generation of MPSoC points to the integration of thousands of IP cores, requiring high performance interconnect for high throughput communications. Optical on-chip interconnect enables significantly increased bandwidth and decreased latency in MPSoC. However, the interface between electrical and photonic devices implies strong layout constraints that may impact the system performance and scalability. In this paper, we propose a novel optical interconnect named CHAMELEON. The interface simplifies the layout and allows the bandwidth between IP cores to be adapted according to the communication requirements. Compared to related networks, CHAMELEON demonstrates improved scalability and flexibility at the cost of minor increase in power consumption.

Keywords—Optical Network on Chip, MPSoC, WDM.

I. INTRODUCTION

Technology scaling down to the ultra deep submicron domain provides for billions of transistors on chip, enabling the integration of hundreds of cores. Many core designs are being increasingly used in modern embedded systems to address the increasing power and performance constraints of embedded applications. Given the increasing number of cores, we are faced with a major challenge in the design of many-core embedded systems: the design and implementation of interconnect that can support high data bandwidth between the cores, as well as between cores and on-chip memories. Designing such systems using traditional electrical interconnect poses a significant challenge: due to capacitive and inductive coupling [12], interconnect noise and propagation delay of global interconnect increase. The increase in propagation delay requires global interconnect to be clocked at a very low rate, which limits the achievable bandwidth and overall system performance. Some attempts were made to solve this problem using different interconnect architectures; however, a new on-chip interconnect technology that can overcome the problems of electrical interconnect is highly desirable.

Optical Network-on-Chip (ONoC) is an emerging technology that is considered as one of the key solutions for the future generation of on-chip interconnects. It relies on optical waveguides to carry optical signals, so as to replace electrical interconnect and provide the low latency and high bandwidth

properties of the optical interconnect. One of the main factors contributing to the impact of an ONoC on the overall performances of an MPSoC is the Optical Network Interface (ONI) architecture. The data rate, the flexibility and scalability as well as the ease of layout synthesis are some examples of MPSoC metrics that are directly related to ONI definition.

In this paper we propose CHAMELEON, which stands for CHANNEL Efficient ONoC, a novel optical interconnect. Compared with existing ONoCs, the main features of the proposed architecture are:

- Higher scalability due to the reuse of layout synthesis offered by a regular ONI and ring topology;
- Highly adaptable bandwidths between IP cores offered by the reconfigurable feature of the ONI;
- Reduced power consumption by assuming the combined use of on-chip lasers and both clockwise (C) and counter-clockwise (CC) directions for signal propagation;
- Higher bandwidth by considering waveguide partitioning to realize independent communications using the same wavelength in the same waveguide.

The paper is structured as follows. Section II discusses the ONoCs proposed in related work. Section III presents CHAMELEON, the proposed ONoC. Section IV gives the evaluation results. Section V concludes the paper.

II. RELATED WORK

ONoCs proposed in related works are divided into two classes, *active* or *passive*, indicating the use of configurable or passive Microring Resonators (MRs) respectively. The type of MRs directly impacts the network scalability and efficiency. Indeed, most *scalable* networks are obtained by using configurable MRs since they allow multiplexed communications in time over photonic devices, leading to circuit switching architectures commonly addressed in electrical Network-on-Chips (NoCs) [2][3]. However, most *efficient* networks rely on passive filters since they do not require any arbitration [1][4] due to the dedicated point-to-point communications. This leads to ORNoC [4], which allows wavelengths to be reused in a same waveguide to design energy-efficient point-to-point channels. CHAMELEON extends ORNoC with a reconfiguration layer to open channels at run-

time, thus allowing better adaptation of the bandwidth according to the application traffic.

Snake [10] is a wavelength-routed optical network providing point-to-point connections between IP cores. A custom place-and-route optimizes the layout by reducing the number of waveguide crossings and the length of waveguide. However, the use of Photonic Switching Elements (PSEs) implicitly requires waveguide crossings that cannot be removed, which significantly impact the optical losses. CHAMELEON does not need any waveguide crossings thanks to the ring topology, which leads to more efficient energy transmission of data.

In ATAC [8], the optical network is used for global broadcasting. Its topology is somewhat similar to CHAMELEON, but the contention-free property is based only on WDM, while in our network it is based on both WDM and wavelength reuse. CHAMELEON has the potential for fewer waveguides/wavelength, which eases scaling to a large-scale architecture. Moreover, contrary to our approach, ATAC does not support simultaneous communications between one source and multiple destinations, unless it is a broadcast of the same message. As opposed to ATAC, our network relies on on-chip laser sources that provide key advantages discussed below.

Efficient on-chip lasers usually require the inclusion of III–V semiconductors: gallium arsenide (GaAs) or indium phosphide (InP) are currently considered to be the best options. Microlasers, based on microdisk structures coupling light evanescently from the cavity resonant mode to the guided mode in an adjacent silicon waveguide, are sufficiently compact as to be implemented in large numbers and at any position. For a given wavelength, the size of an on-chip laser is of the same order of magnitude as the size of a MR used to modulate continuous waves emitted by off chip lasers, which leads to a similar on-chip size for both approaches. While on-chip laser sources require the use of less mature technologies compared to their off-chip counterpart, they have the potential to provide the following three key advantages:

- Easier and more efficient integration by relaxing constraints on layout: it is not necessary to distribute the light from an external source to the modulators (e.g. through the so called *power waveguide* in Corona [7]). Relaxing such constraints contributes to reducing the number of waveguide crossings or even to avoiding them altogether in the ring topology.
- Higher scalability by keeping the architecture fully distributed, which is not achievable by considering centralized off-chip lasers.
- Lower power by reducing the worst case communication distance: source IP $\rightarrow$ destination IP with on-chip laser versus off-chip laser $\rightarrow$ source IP $\rightarrow$ destination IP. This contributes to reducing the propagation losses and consequently the minimum required laser output power. The power consumption can be further improved by locally turning off the laser when no communication is required.

To the best of our knowledge, the network presented in this paper is the first taking advantage of the three above-mentioned key advantages. No waveguide crossing is required and the layout is regular, which make the network implicitly scalable without any custom place-and-route tool [6][10]. These features are achieved thanks to both the use of on-chip lasers and the special architecture design. For example, these

benefits are not observable in architectures like Snake [10], λ -router [1] or WANOc [11] even employing on-chip lasers. CHAMELEON remains fully distributed (including the laser sources) which further contributes to its scalability and also offers the potential for custom design and local run-time control of optical resources, e.g. to turn ON-OFF lasers.

III. CHAMELEON ARCHITECTURE

In this section, we describe CHAMELEON and the ONI. Possible communication schemes are highlighted and an optical loss model is proposed.

A. Architecture Overview

Figure 1 illustrates the considered 3D architecture. It is composed of an electrical layer implementing IP cores and an optical layer implementing CHAMELEON. The optical network in the optical layer is composed of Microring-Resonators (MR) and on-chip laser sources gathered in ONIs. Each of the ONIs is regular and connected to a given IP core. The ONIs are crossed by waveguides propagating optical signals and they are configured at run-time to open dedicated (i.e. point-to-point) communication channels between IP cores. The configuration process is managed by a control network implemented on the electrical layer and is triggered when IP core to IP core communications occur.

Figure 1: CHAMELEON is implemented on the optical layer and it interconnects IP cores located on the electrical layer

The main feature of the network is that it has a regular architecture. Moreover, the ONI architecture is regular and can be reconfigured during run-time. The regularity of the ONI architecture facilitates reuse and layout synthesis, while run-time reconfiguration facilitates low power communications. Moreover, as in ORNoC [4], CHAMELEON allows the reuse of wavelengths to realize several independent communications on a single waveguide. Run-time reconfiguration and wavelength reuse allows the use of the available bandwidth to be adapted according to the communication traffic. The combined use of WDM and multiple waveguides leads to a high overall bandwidth in the optical network.

B. Network interface

The ONI, illustrated in Figure 2, facilitates the communication between IP cores through the optical interconnect. On the technological side, an ONI is composed of an optical part used to propagate data between IP cores through the optical network and an electrical part responsible for the resource allocation through the control network.

Figure 2: Optical Network Interface

Each ONI consists of a receiver part and a transmitter part crossed by a waveguide. The receiver part is composed of wavelength-specific MRs that can be turned ON or turned OFF, in order to respectively configure drop (receive) and pass through operations on the signals at the corresponding wavelength. Signals dropped from the waveguides reach a photodetector, where opto-electronic data conversion generates an electrical signal suitable for electrical receiver circuit. Optical signals that pass through an ONI continue the propagation in the waveguide until they reach the receiver, i.e. an ONI with the MR of the same wavelength that is in the ON state.

The transmitter is composed of on-chip laser sources that can emit and inject optical signals at a specific wavelength into the waveguide. The data are directly transmitted from these lasers through current modulation and each laser source can also be turned OFF in case it is unused.

The receiver and transmitter parts have a symmetric structure: the receiver can eject signals at a given set of wavelengths while the transmitter can inject signals with the same set of wavelengths. In CHAMELEON implementing WDM with N wavelengths, N MRs and N on-chip lasers are used for each waveguide. Each MR or laser source can be configured (i.e. turned ON/OFF) independently. The configuration of these resources is performed by the electrical part of the ONI (i.e. the control network) and it allows the following operations to be realized:

- **injection:** as with sending, the electrical data coming from the IP cores are converted into current, used to control an on-chip laser. For this purpose, the laser must be turned ON. The light is emitted and injected into the waveguide, and then propagates until reaching the receiver part of the destination ONI. Since each laser is wavelength-specific, the selection of the lasers to be used to realize a communication relies on a communication protocol not detailed in this paper. Figure 2 illustrates the injection of signals with wavelength λ_1 (shown in green) into the waveguide.
- **ejection (drop):** optical signal propagating along the waveguide will cross the MRs of the receiver. Signals, whose wavelength matches with the wavelength of MRs in ON state, will be dropped into the perpendicular waveguide, and reach the photodetector (this happens at the receiver part). Figure 2 illustrates the ejection of signals with wavelengths λ_0 and λ_1 from the waveguide (shown in red and green, respectively). It is important to notice that for the remaining part, further along the waveguide, the ejected wavelengths (λ_1 and λ_2) are unused. This allows the reuse of the same wavelength in the

injection part of the ONI in order to realize another communication, as illustrated with the optical signals at λ_1 .

- **pass through:** an optical signal propagating along the waveguide will not be ejected, and no optical signal at the same wavelength is injected for the sake of coherency and to avoid interference. The signal thus crosses the ONI without being modified, as represented by the signal at wavelength λ_2 (blue color) in Figure 2, meaning that the receiver for the signal is further along the waveguide.

When no communication occurs, all the MRs and laser sources are turned OFF for energy saving. They are turned ON/OFF according to the configuration specified by the electrical control network, in order to allocate resources dynamically according to the communication to be realized.

C. Communication Schemes

The configurability of CHAMELEON allows multiple communication schemes to be realized by opening dedicated channels between IP cores, as illustrated in Figure 3:

- **Opening dedicated point-to-point communication channels** is facilitated by waveguide partitioning, which allows ONIs to reuse a given wavelength to realize multiple independent communications in the same waveguide. In Figure 3 a), λ_0 is used to realize communications $ONI_A \rightarrow ONI_B$, $ONI_B \rightarrow ONI_C$ and $ONI_C \rightarrow ONI_A$. Concurrently, λ_1 and λ_2 are used to realize $ONI_C \rightarrow ONI_B$ and $ONI_A \rightarrow ONI_D$ respectively. This facilitates the virtual partitioning of a waveguide for a given wavelength.
- **Broadcast (resp. multicast)** can be realized by opening dedicated communication channels between a source ONI and all the remaining ONIs (resp. the destination ONIs). In Figure 3 b), ONI_B broadcasts data to ONI_C , ONI_D and ONI_A through λ_0 , λ_1 and λ_2 respectively.
- Allocating multiple wavelengths for a given communication can open **high-bandwidth channels**. This is suitable for the execution of streaming applications that temporarily require the transfer of a large amount of data from one IP core to another. In Figure 3 c), high bandwidth communication channels are opened from ONI_B to ONI_D and from ONI_D to ONI_A .
- Multiple waveguides can be used to propagate optical signals in both **clockwise (C)** and **counter-clockwise (CC)** directions. In addition to reducing the worst-case losses in the network (and consequently the power consumption, as will be discussed later on), this allows bi-directional dedicated communication channels to be opened, which will be suitable for processor-memory communications.

These communication schemes can be combined as long as sufficient bandwidth in the network is available. For instance, high bandwidth channels can be opened, while other lower bandwidth channels are already open. This high flexibility makes CHAMELEON suitable to execute applications from various classes (or domains). However, opening channels at the granularity of the wavelength leads to a higher complexity in the control network, which may result in additional latency during the allocation of optical resources to channels. To make CHAMELEON efficient, each channel should thus transmit the largest set of data possible before closing. This suits the streaming model of computation particularly well, since it

usually requires the transfer of large amounts of data for a short period.

Figure 3: Possible communication schemes: a) dedicated distributed point to point communications channel, b) broadcast and c) high bandwidth channel

D. Optical Loss Model

CHAMELEON is composed of on-chip laser sources, MRs, photodetectors and waveguides that introduce optical losses as the signal propagates. In order to evaluate the minimum laser output power $P_{\min_laser}$ (in dBm), the minimum optical power received by the detector $P_{\min_receiver}$ (in dBm) and the worst-case losses in the optical path L_{wc} (in dB) are considered as follows:

$$P_{\min_laser}^{dBm} = L_{wc}^{dB} + P_{\min_receiver}^{dBm}$$

L_{wc} depends on the propagation loss in the waveguide $L_{waveguide}$ (in dB), the through loss $L_{through}$ (in dB) and P_{drop} (in dB), which corresponds to the drop loss occurring when a MR is in the ON state (in CHAMELEON, an optical signal crosses a single MR in drop mode; this occurs in the destination ONI to eject a signal from the waveguide). In addition, we assume that there is negligible bending loss and no waveguide crossing, due to the topology and layout properties [4], which leads to the following equation:

$$L_{wc}^{dB} = L_{waveguide}^{dB} + L_{through}^{dB} + P_{drop}^{dB}$$

$L_{waveguide}$ is obtained from the intrinsic propagation losses of the optical signal in the waveguide $P_{propagation}$ (in dB/cm) and from d_{max} (in cm), the longest distance between the source and destination assuming a serpentine layout. By considering only the C direction for the signal propagation in a network including $N \times M$ ONIs, d_{max} is defined as follows:

$$\begin{cases} d_{max} = (N \times M - 2) \times d + (N + M - 2) \times d & M: \text{odd}; \\ d_{max} = (N \times M - 2) \times d + (M - 1) \times d & M: \text{even}, \end{cases}$$

where d is the distance between two neighboring ONIs.

By considering C-CC (i.e. C and CC directions for signal propagation through the use of separated waveguides), d_{max} is defined as follows:

$$\begin{cases} d_{max} = (\lfloor (N \times M) / 2 \rfloor - 1) \times d + (N + M - 2) \times d & M: \text{odd}; \\ d_{max} = (\lfloor (N \times M) / 2 \rfloor - 1) \times d + (M - 1) \times d & M: \text{even}. \end{cases}$$

$L_{through}$ is the product result between the loss for each MR in through mode $P_{through}$ (in dB) and $N_{through}$, the maximum number of MRs in the through mode passed by an optical signal at the corresponding wavelength. By considering the C direction, $N_{through}$ equals $(N \times M - 2)$. By considering C-CC directions, $N_{through}$ equals $\lfloor (N \times M) / 2 \rfloor$ or

$(N \times M - 1 - \lfloor (N \times M) / 2 \rfloor)$ when M is odd or even, respectively.

It is worth noticing that d_{max} and $N_{through}$ are significantly reduced for the C-CC case compared to the C case. This will result in a lower worst-case loss, which directly contributes to the energy-efficiency of the network. Indeed, for a given photodetector responsivity and a given target Bit Error Rate (BER), the lower loss in the communication path results in a lower minimum required laser output power.

IV. RESULTS

We compare CHAMELEON with Snake [10], ORNoC [4] and SWMR (Single Write Multiple Read), which is modeled on ATAC [8]. Snake and ORNoC are passive networks relying on multistage and ring topologies respectively. For CHAMELEON and ORNoC, we consider both C-only and C-CC directions for signal propagation, thus leading to CHAMELEON_C, CHAMELEON_{C-CC}, ORNoC_C and ORNoC_{C-CC}.

A. Architectures

The comparison is achieved by considering 3 architectures:

Arch₁ is extracted from [10] and is a processor to memory application. Figure 4 a) illustrates the layout considered for Snake: it is adapted from [10] to match the requirements of a fully-integrated system in which 4 processors ($P_0 \dots P_3$) and 4 memories ($M_0 \dots M_3$) share the same 20x10mm² electrical layer. Processors are interconnected through a crossbar located in the center (not shown in the figure), which avoids placing Snake in this area. We assume a 5mm distance d between optical interfaces in the optical network. The layouts for ORNoC and CHAMELEON involve closed waveguides successively crossing $M_0, M_1, P_1, P_3, M_3, M_2, P_2$ and P_0 . For a fair comparison with CHAMELEON, we assume that on-chip lasers are used in Snake.

Arch₂ corresponds to 4x4 IP cores ($IP_0 \dots IP_{15}$) connected with the optical network. A 20x20mm² die size is assumed and $d=5$ mm. Figure 4 b) represents the layout for Snake which is designed to avoid any waveguide crossing between the network interfaces and the Snake multistage itself (Snake is located in the middle of the optical layer for layout optimization purposes and is represented as a box for the sake of clarity). Snake interconnects 16 inputs (in red lines) with 16 outputs (in black lines) through 112 PSEs. The initial structure

of Snake would assume 120 PSEs but, for a fair comparison, we adapt the reduction method from [1] to Snake in order to remove unused PSEs. The layout we assume for CHAMELEON and ORNoC is the one illustrated in Figure 1.

Arch₃ extends *Arch₂* to 8x8 IP cores, thus matching the ATAC architecture: 20x20mm² die size and $d=2.5$ mm are assumed. The size of Snake is increased to match the new connectivity requirement, and the layouts of CHAMELEON and ORNoC are extended from *Arch₂*.

Table 1 summarizes the parameters we used to compare the networks. Similarly to [8], we consider conservative (Co) and aggressive (Ag) values.

Table 1: Injection Loss parameters

Optical loss	Conservative (Co)	Aggressive (Ag)
$P_{\text{propagation}}$	1.5 dB/cm [10]	0.2 dB/cm [8]
P_{drop}	0.013 dB [10]	1 dB [8]
P_{through}	0.05 dB [6]	0.01 dB [8]
P_{crossing} (Crossing loss)	0.52 dB [10]	0.05 dB [6]

Figure 4: Considered Snake [10] layout for a) *Arch₁* and b) *Arch₂*

B. Network Comparisons

We evaluate the following ONoC characteristics: worst-case optical loss (L_{wc}) considering both Ag and Co values, number of waveguides (N_{wg}), number of wavelengths per waveguide (N_{wl}), number of on-chip lasers (N_{laser}), and number of MRs (N_{MR}) (NB: for Snake, the number of MRs takes into account both the MRs based filters in the receiver part of the network interface and the MRs based PSEs in the network itself - one PSE counts for 2 MRs).

Regarding *Arch₁*, we estimate the worst case distance for Snake as follows: from P_2 to M_2 , the signal will propagate through a distance estimated to be equivalent to 4 times the minimum distance between ONIs (i.e. d). Since this worst-case path also suffers from 6 waveguide crossings (4 through the PSEs), the total loss is estimated to be 1.7dB and 6.133dB for Ag and Co values respectively (considering multi-layer silicon deposited technology, which allows 3D photonic devices, would contribute to reducing the losses [6]). Snake is composed of 56 MRs, including 12 PSEs (as indicated in the bracket in Table 2). It requires 4 wavelengths per waveguide, which is the limit considered for ORNoCs and CHAMELEONS. ORNoC_C shares similar characteristics to Snake when considering conservative values. However, because ORNoC does not suffer from any waveguide crossings, a significant improvement is obtained when considering aggressive values (1.7dB and 0.7dB for Snake and ORNoC_C respectively). Further improvements are obtained with ORNoC_{C-CC} since the CC direction for signal propagation allows the length to be

reduced from 35mm to 20mm. CHAMELEON directly inherits from the main features of ORNoC: no waveguide crossings, and the possibility to combine C and CC rotations. Both features allow the number of waveguides to be reduced from 4 (CHAMELEON_C) to 2 (CHAMELEON_{C-CC}) and the number of on-chip lasers from 128 to 64 respectively. The extra MRs located in the receiver part of the ONIs introduce through losses (0.03dB) in the worst-case path for CHAMELEON_{C-CC}. The overhead in the number of on-chip lasers in CHAMELEON is due to its reconfigurability property. Indeed, the complexity of CHAMELEON (e.g. number of lasers) is defined in order to allow the same connectivity in Snake and ORNoC to be configured for the considered example architecture. Moreover, CHAMELEON can open/close communication channels at run-time. This allows, for instance, additional bandwidth to be allocated for a given memory to processor channel by closing other channels. As another example, new channels can be opened between 2 processors, which is impossible for Snake and ORNoC unless it is specified at design time. Since unused lasers are turned-off, CHAMELEON does not suffer from extra power consumption. As a primary conclusion, CHAMELEON offers a run-time flexibility to adapt the bandwidth distribution according to the connectivity requirements at the price of acceptable extra losses compared to ORNoC (the best solution mentioned in [10]).

Table 2: Comparisons of CHAMELEON with related ONoCs

	Snake	ORNoC _C	ORNoC _{C-CC}	CHAMELEON _C	CHAMELEON _{C-CC}
<i>Arch₁</i>	N_{laser}	32	32	128	64
	N_{wl}	4	4	4	4
	N_{wg}	8	4	2	4
	N_{MR}	56 (12)	32	32	128
	$L_{wc} Ag$	1.7	0.7	0.4	0.76
	$L_{wc} Co$	6.133	5.25	3	5.55
<i>Arch₂</i>	N_{laser}	240	240	1,920	960
	N_{wl}	15	15	15	15
	N_{wg}	16	8	4	8
	N_{MR}	464 (112)	240	240	1,920
	$L_{wc} Ag$	2.9	1.7	1	1.84
	$L_{wc} Co$	16.8	12.75	7.5	13.45
<i>Arch₃</i>	N_{laser}	4032	4,032	4,032	129,024
	N_{wl}	63	63	63	63
	N_{wg}	64	32	16	32
	N_{MR}	8,000 (1,984)	4,032	4,032	129,024
	$L_{wc} Ag$	5.5	3.45	1.9	4.07
	$L_{wc} Co$	43.26	25.875	14.25	28.975

Arch₂: Snake and ORNoC are crossbars. CHAMELEON basically follows the same trend: close to ORNoC but with more flexibility. By assuming 1GHz modulation speed for the lasers, CHAMELEON offers the same bandwidth as ORNoC and Snake (which can be estimated at 240 Gbit/s, i.e. 240×1Gbit/s) when configured as a crossbar allocating one wavelength per channel between IP cores. However, if we consider the execution of a streaming application where data propagate from an IP to another, CHAMELEON has the potential to deliver 1.92 Tbit/s bandwidth by turning ON all the laser sources.

Arch₃ highlights the resources overhead of CHAMELEON when designed to allow the configuration of a full crossbar between 64 IP cores. For such a large-scale system, permanent connectivity between all the IP cores may not always be required, which could justify a reduction of the number of lasers in CHAMELEON and, therefore, a reduction of its flexibility. Design tradeoffs thus need to be explored by simulating the execution of representative benchmarks. SWMR only requires 64 off-chip lasers to implement a broadcast, which may lead to less efficient energy/bit transmission.

Finally, for Co values, worst case loss for SWMR is 16.06dB compared to 14.25dB and 15.8dB for ORNoC_{C-CC} and CHAMELEON_{C-CC} respectively.

C. Power Efficiency of CHAMELEON

We evaluate the total laser output power required for the evaluated worst-case loss in *Arch3* and for the target 10^{-12} BER [6]. Since we consider a germanium photodetector with the responsivity of 1A/W, the minimum received power is consequently -20dBm (10 μ W) for error-free operation with the target BER. Figure 5 represents the estimation for the aggressive values. The smaller waveguide for Snake does not alleviate the high number of waveguide crossings specific to the multistage topologies (24.45mW and 332W for Ag and Co respectively). ORNoC_{C-CC} is the most power efficient network, requiring 10.67mW (Ag) and 417mW (Co). Compared to ORNoC_{C-CC}, CHAMELEON_{C-CC} requires 7.4% and 42% additional power for Ag and Co respectively, which appears acceptable considering its reconfigurability features. However, simulations are required to evaluate its run-time behavior since, on one hand, we assumed the network already configured as a crossbar (thus not considering the reconfiguration time) and, on the other hand, we do not take advantage of the potential of CHAMELEON to reduce the system power or to improve the execution performances by adapting the bandwidth to the application traffics.

Figure 5: Power efficiency of CHAMELEON

V. CONCLUSION

In this paper, we presented CHAMELEON, a novel ONoC that makes good use of WDM to create communication channels between IP cores. To the best of our knowledge, this network is the first allowing the run-time creation of point-to-point (i.e. dedicated) channels without any waveguide crossing in the optical path, which leads to energy-efficient optical transmission of data. Compared to related static (i.e. non-configurable) ONoCs designed to fully interconnect 8x8 cores, CHAMELEON can be configured at run-time to realize the same connectivity with an energy overhead of 7.4% when compared to the most energy-efficient non-configurable solution. The ring topology and the regular layout of the interfaces contribute

to the good scalability of CHAMELEON. The combined use of clockwise and counter-clockwise directions for signal propagation allows a substantial improvement of its energy-efficiency and scalability. The reconfigurable ability of CHAMELEON allows the bandwidth to be adapted between IP cores according to application traffic requirements, which will further reduce the energy/bit transmission of data for a given application. This will be further evaluated in future works.

ACKNOWLEDGMENT

Sébastien Le Beux is supported by a Région Rhône-Alpes grant.

REFERENCES

- [1] I. O'Connor, et al. "Reduction Methods for Adapting Optical Network on Chip Topologies to Specific Routing Applications". In Proceedings of DCIS, November 2008.
- [2] A. Shacham, K. Bergman, L.P. Carloni, "Photonic Networks-on-Chip for Future Generations of Chip Multi-Processors," IEEE Transactions on Computers 57 (9), pp. 1246-1260, 2008
- [3] Y. Ye et al. "A Torus-based Hierarchical Optical-Electronic Network-on-Chip for Multiprocessor System-on-Chip", ACM Journal on Emerging Technologies in Computing Systems, 2012.
- [4] S. Le Beux, et al. "Layout Guidelines for 3D Architectures including Optical Ring Network-on-Chip (ORNoC)". In 19th IFIP/IEEE VLSI-SOC International Conference, 2011
- [5] 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". Proceedings of the Third International Symposium on Networks-on-Chip (NOCS), 2012.
- [6] A. Biberman, K. Preston, G. Hendry, N. Sherwood-Droz, J. Chan, J. S. Levy, M. Lipson, 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 7 (2) 7:1-7:25
- [7] D. Vantrease, et al. Corona: System Implications of Emerging Nanophotonic Technology. In Proceedings of the 35th Annual International Symposium on Computer Architecture (ISCA) pages 153–164, 2008.
- [8] J. Psota, et al. 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.
- [9] J. Van Campenhout et al., "A compact SOI-integrated multiwavelength laser source based on cascaded InP microdisks," IEEE Photon. Technol. Lett., vol. 20, no. 16, pp. 1345–1347, 2008
- [10] Luca Ramini, Paolo Grani, Sandro Bartolini, and Davide Bertozzi. Contrasting wavelength-routed optical NoC topologies for power-efficient 3D-stacked multicore processors using physical-layer analysis. In Proceedings of the Conference on Design, Automation and Test in Europe (DATE), 2013.
- [11] Zheng Chen, Huaxi Gu, Yintang Yang and Ke Chen. Low Latency and Energy Efficient Optical Network-on-Chip Using Wavelength Assignment. In IEEE Photonics Technology Letters, Vol. 24, Issue 24, 2012
- [12] R. Ho, K.W. Mai, and M.A. Horowitz. The Future of Wires. Proceedings of the IEEE, 89(4):490–504, April 2001