Reducing power consumption of lasers in photonic NoCs through application-specific mapping

EDOARDO FUSELLA and ALESSANDRO CILARDO, University of Naples Federico II

To face the complex communication problems that arise as the number of on-chip components grows up, photonic networks-on-chip have been recently proposed to replace electronic interconnects. However, photonic networks-on-chip lack efficient laser sources, possibly resulting in an inefficient or inoperable architecture. In this paper, we introduce a methodology for the design space exploration of optical NoC mapping solutions, which automatically assigns IPs/cores to the network tiles such that the laser power consumption is minimized. The experimental evaluation shows average reductions of 34.7% and 27.3% in the power consumption compared to respectively application-oblivious and randomly mapped photonic NoCs, allowing improved energy efficiency.

(pre-print)

CCS Concepts: • **Hardware → Emerging optical and photonic technologies**; *Photonic and optical interconnect*; *Network on chip*;

Additional Key Words and Phrases: Silicon Photonics, Optical Network-on-Chip, Design automation, On-chip interconnects, Application mapping, Laser, Power consumption

ACM Reference format:
Edoardo Fusella and Alessandro Cilardo. 2018. Reducing power consumption of lasers in photonic NoCs through application-specific mapping. *ACM J. Emerg. Technol. Comput. Syst.* 1, 1, Article 1 (January 2018), 11 pages.
https://doi.org/10.1145/3173463

1 INTRODUCTION

The continuous growth in application performance requirements has drastically changed the scale of multiprocessor systems-on-chip (MPSoCs). Current highly parallel MPSoCs consist of tens to hundreds of cores on a single die, requiring a high-bandwidth, low-latency and energy-efficient network-on-chip (NoC). However, as the network scales up, traditional electronic interconnects fail in fulfilling these requirements [3]: at the deep submicron scale, metallic interconnects are susceptible to non-negligible parasitic resistance and capacitance resulting in poor performance and energy efficiency.

Silicon photonics [19] has generated an increasing interest over the last few years for optical interconnects in integrated circuits, providing a promising answer to effectively face the power

This work has been partially supported by the EC H2020 MANGO project (Agreement No. 671668).
Author's addresses: Department of Electrical Engineering and Information Technologies, University of Naples Federico II, via Claudio 21, 80125 Naples, Italy; emails: {edoardo.fusella, acilardo}@unina.it.

1550-4832/2018/1-ART1 $15.00
https://doi.org/10.1145/3173463

wall, today seriously limiting further technology advances. In particular, nanophotonic waveguides can achieve bandwidths in the order of terabits per second by exploiting wavelength division multiplexing (WDM), while photonic signaling is expected to consume less power than electrical interconnects [6].

Several MPSoCs exploiting photonic NoCs have been proposed [9, 14, 22, 26, 28]. However, silicon photonics lacks efficient native-substrate laser sources to drive the photonic links. Solutions based on both off-chip and integrated on-chip sources have been proposed in the literature. Off-chip laser sources introduce high coupling losses, low flexibility and packaging issues as well as poor energy proportionality making the on-chip counterpart [11] preferable. On the other hand, one of the main drawbacks of on-chip laser sources is the low wall-plug efficiency, due to the conversion of the electrical power into optical power, which depends on both the required output power and the temperature of the laser source, involving power-hungry techniques for cooling and temperature stabilization.

Some works proposed novel III-V semiconductor lasers that are heterogeneously integrated on-chip with the CMOS devices [24] or use metamorphic [27] and pseudomorphic [23] growth techniques. Although these approaches require further development, they exhibit promising improvements. Further enhancements can be achieved through optimizations at the architectural level. In that respect, Kurian *et al.* [13] highlighted the importance of having on-chip lasers that allow rapid power gating in order to switch on and off the laser devices in an efficient way. Chen *et al.* [2] proposed to share each laser source across different power waveguides in order to enable the lasers to work at their peak efficiency. In addition, the placement of lasers is evaluated and optimized in such a way that they can operate at the minimum temperature as possible. Ye *et al.* [30] presented a torus-based optical NoC exploiting an adaptive power control technique to properly estimate the adequate laser output power required for each path.

From a different perspective, in the design automation community, a typical design flow consists of the following steps. First, the application is partitioned into a set of concurrent tasks. Second, the application tasks are assigned and scheduled into a given set of available intellectual property (IP) blocks. These IPs range from CPUs, DSPs, customized application-specific integrated circuits (ASICs), embedded DRAMs and the like. Finally, IPs are topologically placed onto the different tiles of a target NoC architecture such that the metrics of interest are optimized (see Fig. 1). The third part of system design flow, which is the application mapping problem, is the focus of our work. While there is a large body of work focusing on the mapping problem for electronic NoC architectures, there are only a few works targeting the photonic counterpart aiming to optimize different metrics of interest [5, 7]. In contrast, as the main contributions of this paper, we first formulate the application-specific mapping optimization problem for photonic NoC architectures with the objective of minimizing the laser energy consumption under bandwidth constraints. Then, we provide and compare three different algorithms to solve it. The experimental evaluation shows that the laser power consumption can be significantly reduced, allowing improved energy efficiency. To the best of our knowledge, our work is the first to propose an application-specific optimization as a way to reduce the laser power consumption in photonic NoC architectures.

2 PRELIMINARIES

In this section, we briefly describe the photonic NoC architecture and the related insertion loss and laser power consumption models. We consider a tile-based implementation laid out as a mesh or torus topology, as shown in Fig. 1. Adjacent nodes are connected by two unidirectional silicon waveguides. In addition, torus topologies are enhanced with wrap-around waveguides between the edge nodes. Each tile contains an IP core and an optical router and is connected with the four

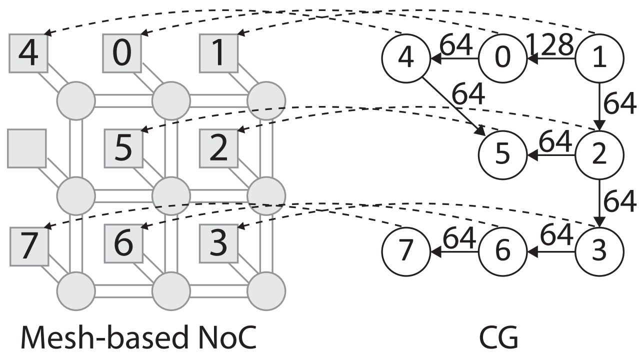

Fig. 1. A mesh-based on-chip architecture and an example of mapping problem.

neighboring tiles. External lasers coupled to on-chip power waveguides distribute light across the chip. Laser sources are placed along the edges of the chip and shared between two or more waveguides to improve the laser efficiency and power consumption [2].

Concerning the routing algorithm, we employ a minimal deterministic dimension order routing (DOR) since it is easy to implement, enforces packets to take only shortest paths and ensures deadlock/livelock freedom. In addition, in case of DOR, optical switches can be designed with straight default paths[1] and without support for those turns that are not allowed by DOR, leading to a lower number of waveguide crossings and MRs, and hence reduced insertion loss.

2.1 The optical loss and laser models

Photonic NoCs are composed of several devices (waveguides, microrings, modulators, etc...) that introduce optical power loss, affecting photonic signals as they propagate along a path. The power loss from a source to a destination $L^{(src,dst)}$ can be evaluated as the sum of all the losses in each hop along the path between the two end-points according to equation 1

$$L^{(src,dst)} = L_{mod} + L_{coup} + L_{top} \tag{1}$$

where

- L_{mod} is the loss due to the electro-optical modulator
- L_{coup} is the loss due to the couplers used to interface with the off-chip components
- $L_{top} = L_{prop} + L_{cross} + L_{bend} + L_{drop} + L_{pass}$ is the sum of all the losses affecting a signal due to topological choices:
  $L_{prop} = P_{prop} \times d$ is the loss affecting a signal when it propagates in a straight waveguide with a length equal to d
  $L_{cross} = P_{cross} \times n_{cross}$ is the loss due to crossing other waveguides
  $L_{bend} = P_{bend} \times n_{bend}$ is due to waveguide bends
  $L_{drop} = P_{drop} \times n_{drop}$ is due to dropping into a ring
  $L_{pass} = P_{pass} \times n_{pass}$ is due to passing by a ring

with P_{prop}, P_{cross}, P_{bend}, P_{drop}, P_{pass} being the unitary loss parameters and n_{cross}, n_{bend}, n_{drop}, n_{pass} the number of occurrences in the path between the two end-points. Table 1 shows some unitary insertion loss parameters. The data, which are a projection toward the optical interconnect technology in 2020, are obtained from [20] and [4]. Please note that, while we rely on this table for

[1] A default path is the path that the signal takes when all the rings are placed in an off resonance state.

the power loss estimation, the proposed approach is independent of the loss coefficients actually used.

Table 1. Loss parameters

Parameter	Notation	Value
Modulator	P_{mod}	-0.6 dB
Coupler	P_{coup}	-0.7 dB
Propagation Loss in Silicon	P_{prop}	-0.274 dB/cm
Crossing loss	P_{cross}	-0.04 dB
Bending loss	P_{bend}	-0.005 dB/90°
Dropping into a Ring	P_{drop}	-0.5 dB
Passing by a Ring	P_{pass}	-0.005 dB

Obviously, the power loss is highly dependent on the optical switch microarchitecture. In that respect, we rely on the Crux switch, first presented in [29]. Crux is designed in a power loss-aware way with straight default paths and optimized for XY routing. In fact, the total number of MRs is only 12, while the total number of waveguide crossings is 9, letting light waves cross at most three crossings and a single MR in the worst-case path.

Notice that, in order to properly translate signals into the electrical domain, received optical waves need to have a minimum power above the photodetector sensitivity. Usually, the worst-case power loss is used to set the laser source in order to provide the worst-case optical power for all the optical signals. In such a case, the power generated by the laser sources must be equal to the sum of the worst-case power loss and the photodetector sensitivity. However, this leads to a power waste for all those communications that are subject to a lower power loss. For this reason, similar to [30], we rely on an adaptive power control technique that uses topology and routing information to evaluate the loss of a certain path and drive the laser in order to generate just enough power for that path. This leads to the following equation:

$$L^{(src,dst)} + S_{dst} = P_{out}^{(src,dst)} \tag{2}$$

where S_{dst} is the photodetector sensitivity of the destination node and $P_{out}^{(src,dst)}$ is the laser output power generated in the node *src* required to reach the node *dst* that can be evaluated according to equation (3) as a function of the required laser input power $P_{in}^{(src,dst)}$ and the wall-plug efficiency η_{WPE}.

$$P_{out}^{(src,dst)} = P_{in}^{(src,dst)} \eta_{WPE} \tag{3}$$

Based on the above equations, it is possible to evaluate the laser power consumption as a function of the photodetector sensitivity, the power loss along the path, and the wall-plug efficiency.

3 METHODOLOGY

3.1 Problem formulation

This work deals with the mapping problem for photonic tile-based NoC architectures. For a given application, our objective is to decide on which tile should each core be mapped such that the laser power consumption is minimized under given bandwidth constraints. To this end, we need the following definitions.

Definition 3.1. A *Communication Graph* is a directed graph $CG = G(C, E)$ with each vertex $c_i \in C$ representing an IP core and the directed edge (c_i, c_j), denoted as $e_{i,j} \in E$, representing the communication between cores c_i and c_j. Each $e_{i,j}$ has an attribute $b(e_{i,j})$ expressing application-specific information, i.e. the communication bandwidth requirement.

Definition 3.2. A *Topology Graph* is a directed graph $TG = G(T, L)$ with each vertex $t_i \in T$ representing a tile in the network and the directed edge (t_i, t_j), denoted as $l_{i,j} \in L$, representing a physical optical link between the tiles t_i and t_j. Each $l_{i,j}$ has an attribute $B(l_{i,j})$ that gives the available bandwidth of the link. Without loss of generality, we assume that all the links have the same bandwidth B, and hence $B(l_{i,j}) = B \quad \forall l_{i,j} \in L$.

Using the above graph representations, the problem addressed can be formulated as follows.
Given a CG and a TG satisfying

$$size(C) \leq size(T) \tag{4}$$

Find a mapping function $\Omega : C \rightarrow T$ which minimizes

$$\min \left\{ P_{laser}^{tot} = \sum_{\forall e_{i,j} \in E} P_{in}^{(\Omega(c_i), \Omega(c_j))} \frac{b(e_{i,j})}{B} \right\} \tag{5}$$

Such that:

$$\forall c_i \in C, \quad \Omega(c_i) \in T \tag{6}$$

$$\forall c_i \neq c_j \in C, \quad \Omega(c_i) \neq \Omega(c_j) \tag{7}$$

$$\forall l_{i,j} \in L, \quad \sum_{\forall e_{i,j} \in E} b(e_{i,j}) \times f(l_{i,j}, e_{i,j}) \leq B \tag{8}$$

where

$$f(l_{i,j}, e_{i,j}) = \begin{cases} 1 & \text{if } e_{i,j} \text{ is routed on the optical link } l_{i,j}, \\ 0 & \text{otherwise.} \end{cases} \tag{9}$$

P_{laser}^{tot} gives the power consumption of all the laser sources in the network. Conditions (6) and (7) guarantee respectively that each core should be mapped to exactly a single tile and no tile can host more than one core. Last, condition (8) specifies the performance constraints in terms of the aggregated bandwidth requirements for each optical link. Finally, note that, based on equations (2) and (3), $P_{in}^{(\Omega(c_i), \Omega(c_j))}$ can be calculated as $L^{(\Omega(c_i), \Omega(c_j))} \frac{S_{dst}}{\eta_{WPE}}$.

3.2 Design Space Exploration

The application mapping problem is a specialization of the constrained quadratic assignment problem which is known to be NP-hard [10]. It was proven that there isn't any algorithm for solving this problem in polynomial time and, hence, it is usually solved using heuristic techniques. In that respect, we rely on three different algorithms: a random search (RS), a genetic algorithm (GA), and a randomized priority-based list algorithm (R-PBLA). The random search simply chooses the best solution among a population of a given size.

The genetic algorithm uses a population of constant size and guides the evolution of a set of selected individuals through a number of generations based on the statistics of the generation. Each phenotype, i.e. candidate solution, has a genotype, i.e. its set of properties, which can be altered using the crossover and mutation operators. Concerning the crossover operator, we rely on the cycle crossover [21], guaranteeing that conditions (6) and (7) are met. Differently, the mutation operator swaps the position of two cores in order to provide a new and feasible solution, thereby increasing the exploration of search space. A fitness value $f_i = 1/P^{tot}_{laser}$ is used to evaluate the solution domain, while for the selection operator we rely on the roulette wheel selection with a probability $p_i = f_i/\sum_{j=1}^{P_{size}} f_j$, where P_{size} is the size of the population.

Finally, in R-PBLA a list of candidate solutions is created first, and then, sorted by laser power consumption. In each generation, the minimum value in the list is used as the candidate solution to calculate the next generation list. The list is made up of all the mapping solutions that can be generated starting from a single candidate solution and considering all the moves that swap the position of two cores. Note that, to avoid ending up with a local optimum solution, when the algorithm reaches such a solution, a new list is generated starting from a new random mapping solution.

4 CASE STUDIES

We validate the effectiveness of the proposed technique using eight real-life applications from the multimedia and networking domains, namely *263dec* (assigned and scheduled onto 14 cores), *263enc* (12 cores), *DVOPD* (32 cores), *MPEG-4* (12 cores), *MWD* (12 cores), *PIP* (8 cores), *VOPD* (16 cores), and *Wavelet* (22 cores). The CGs were obtained from [25] and their characteristics are summarized in Table 2, including the minimum NoC size required to host the applications.

Table 2. Applications characteristics

Application	$\|C\|$	$\|T\|$	**Bandwidth (Mb/s)**	**NoC size**
263dec	14	15	1.37	4×4
263enc	12	12	19.18	3×4
DVOPD	32	44	199.14	6×6
MPEG-4	12	26	256.74	3×4
MWD	12	11	93.33	3×4
PIP	8	8	72	3×3
VOPD	16	21	185.75	4×4
Wavelet	22	35	80.28	5×5

The proposed methodology was implemented and embedded in *PhoNoCMap* [8], a tool for the design space exploration of optical NoC mapping solutions. We assume a $400mm^2$ die area (A_{die}) and we compute the waveguide lengths for an $M \times N$ network as $\sqrt{A_{die}/((M-1) \times (N-1))}$. We consider 20 wavelength channels [15], and a 10 Gbps fixed modulation rate per wavelength [1]. The wall-plug efficiency is set to 10%, the maximum achievable value in case of an on-chip laser device taking into account its temperature, laser source length, and required optical power per wavelength [2]. Last, all nodes employ photodetectors with a sensitivity of −14.2 dBm and a data rate of 10 Gbps, as demonstrated in [17]. Table 3 summarizes the used architectural parameters.

Table 3. Architectural parameters

Parameter	**Value**
Chip size (mm^2)	20×20
Modulation rate (Gb/s)	10
# Wavelength Channels	20
Wall-plug efficiency	10%
Photodetector sensitivity (dBm)	-14.2

We first carried out a set of experiments to evaluate the impact of different mapping solutions on the laser power consumption. For each application, we generated randomly 100, 000 mapping solutions targeting a mesh-based photonic NoC and we evaluated the laser power consumption related to each mapping solution. The results in Fig. 2 show, for each application, the probability distribution of the laser power consumption according to the different mapping solutions. It can be easily recognized the high variability of laser power consumption: although the optimal solution is not necessarily included, on the average, the worst solution requires approximately 36% more power compared to the best randomly generated solution. Note that applications with higher bandwidth requirements involve higher laser power consumption. This is because, in the absence of communication on the optical link, lasers are turned off.

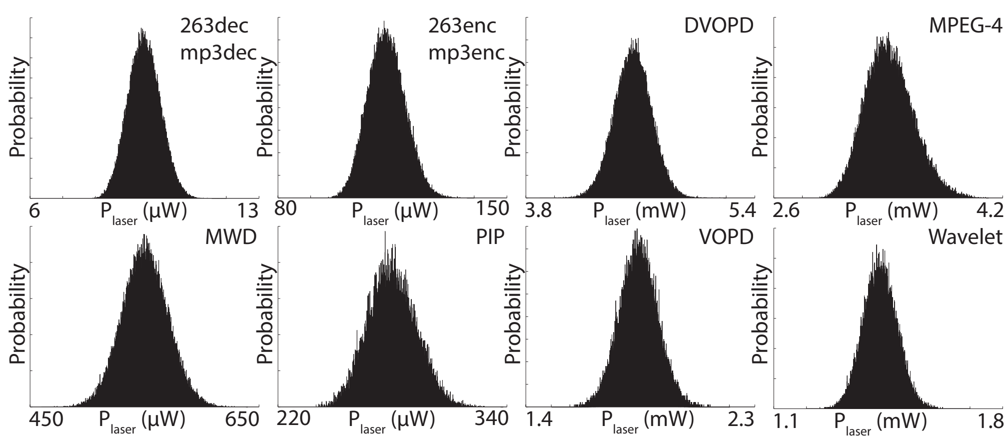

Fig. 2. Probability distributions of the laser power consumptions related to 100, 000 mapping solutions randomly generated for the eight applications.

In a second set of experiments, we compared the best mapping solution found with our R-PBLA algorithm with mesh and torus-based photonic NoCs without a application-specific mapping optimization. In such a case, in absence of any information on the mapping, lasers should provide the worst-case optical power for all the optical signals. Fig. 3 depicts the laser power consumption improvement. On the average, lasers of photonic NoCs exploiting our approach consume 34.7% less power compared to the non-optimized counterpart.

In a third set of experiments, the different mapping algorithms were used under the same temporal bound to find the best solutions. Table 4 shows, for each application, the laser power consumption P_{laser}^{tot} (mW). The best solutions found with the RS strategy (under the given search time constraint) are, on the average, 17% more efficient in terms of laser power consumption compared to the average value over the 100,000 randomly generated solutions with no mapping optimization. Further improvements are obtained with GA and R-PBLA: on average, GA outperforms RS by 5.15%, while R-PBLA finds mapping solutions leading to a further 5.18% laser power consumption reduction compared to GA.

Then, we compared the runtime of the different algorithms. First, we found the best mapping solution by running R-PBLA for 1 minute. Then, we used RS and GA algorithms to find a solution with the same cost of the solution found with R-PBLA. The results are depicted in Table 5. In case of large applications, i.e. DVOPD and Wavelet, both RS and GA were not able to find an appropriate solution after two hours of optimization. This also happens with the RS algorithm and the VOPD

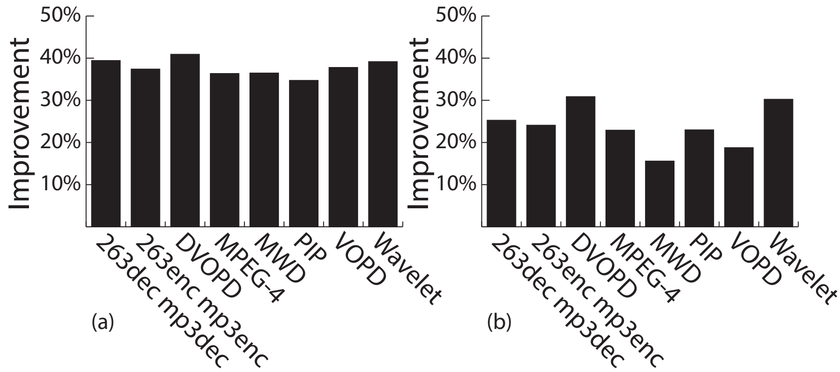

Fig. 3. A comparison between an optical NoC without an optimization mapping strategy and an optical NoC with a mapping found with our R-PBLA algorithm for the following topologies: (a) mesh and (b) unfolded torus.

Table 4. Algorithms comparisons

Application	Mesh			Torus		
	RS	**GA**	**R-PBLA**	**RS**	**GA**	**R-PBLA**
263dec	8.83	8.11	7.89	9.56	9.08	8.86
263enc	114.66	108.42	95.42	116.66	109.67	106.12
DVOPD	4524.31	4489.21	4009.23	4996.23	4857.61	4693.38
MPEG-4	3262.38	3087.20	2921.56	3347.35	3262.38	3245.11
MWD	557.92	552.99	471.20	625.89	577.41	574.12
PIP	278.873	267.61	239.97	289.83	277.63	268.49
VOPD	1889.18	1789.50	1573.18	1973.69	1721.18	1869.76
Wavelet	1455.12	1400.35	1241.10	1585.14	1463.35	1480.77

application. In general, R-PBLA is over an order of magnitude faster than the other algorithms. Only, in case of very small applications, i.e. PIP, the other algorithms outperform R-PBLA.

Table 5. Runtime of the proposed mapping algorithms

Application	RS	GA
263dec	1.07 h	53.02 min
263enc	1.8 min	5.62 min
DVOPD	+ 2h	+ 2h
MPEG-4	20.89 min	17.92 min
MWD	9.82 min	8.51 min
PIP	7.28 s	29.25 s
VOPD	+ 2h	1.6 h
Wavelet	+ 2h	+ 2h

Moreover, note that mapping optimization techniques may have a possible impact on the performance of the whole network. Silicon photonics can benefit from very low latencies since, with no conflicts, the propagation delay is simply the time of flight at the speed of light. As a consequence, in absence of congestion, the latency in the photonic domain does not depend on the different mapping solutions. However, optical NoCs based on regular direct topologies are often supported by an electronic layer where an electronic NoC acts to establish the end-to-end optical paths. We, thus, compared the solutions found with the proposed approach with 100.000 randomly generated mapping solutions in terms of hop count targeting a mesh-based photonic NoC. Although the hop count does not take into account complex effects of the network, such as congestion, it is used in similar works [12, 18] to estimate the performance of electronic on-chip networks. Table 6 depicts the results. On average, our approach leads to a hop count reduced by 27% compared to the average count of the random mappings, with 7.5% more hops compared to the best mapping solution among the 100.000 randomly generated solutions.

Table 6. Hop count comparison

Application	Random (Average)	Random (Best)	Random (Worst)	Proposed
263dec	3.335	2.417	4.929	2.612
263enc	3.666	2.571	4.417	2.783
DVOPD	4.997	3.864	5.955	4.124
MPEG-4	3.331	2.462	4.385	2.659
MWD	3.337	2.333	4.250	2.635
PIP	3.666	2.571	4.929	2.675
VOPD	3.666	2.600	4.500	2.750
Wavelet	4.335	3.371	5.343	3.771

Finally, we present a scalability analysis. We consider four real applications included in the MCSL [16] realistic NoC traffic suite, namely: *FFT-1024_complex*, *Fpppp*, *H264-1080p_dec* and *Sparse*. The applications have been partitioned into a set of tasks that are mapped to regular NoCs with four different sizes: 4×4, 6×6, 8×8 and 10×10. Table 7 shows the laser power consumption P_{laser}^{tot} (mW) related to the solutions found with R-PBLA for the four network sizes. The differences between the different applications can be explained taking into account the CG completeness index (CGCI)[2], shown in Table 8. Generally, having a lower value leads to better results. This is because this index summarizes how much the design space is constrained by the number of communications. For instance, application Sparse achieves the lower laser power consumption regardless of the network size since the application has the smallest CG completeness index. Similarly, application FFT-1024_complex, characterized by the higher CGCI, achieves the worst power consumption values. In addition, Table 7 gives the run times of R-PBLA to find the mapping solutions. In the worst case, i.e. application Sparse mapped to a 10×10 network, it requires 15 minutes and 7 seconds. Note that this is an appreciable result considering that the solution space grows factorially with the network size and hence the are 100! different solutions in a 10×10 network.

5 CONCLUSIONS

Photonic on-chip communication is today considered a major pathway to energy-efficient ultra-high bandwidth on-chip communication. However, silicon photonics lacks efficient laser sources, which leads to the use of power-hungry laser devices that are hard to integrate into future commercial systems. In this paper, we introduced a methodology which aims to reduce the laser power

[2]The CG completeness index has been derived as the ratio between the number of arcs in the CG and $size(C) \cdot (size(C) - 1)$, the maximum possible number of arcs in absence of self-loops.

Table 7. Comparison under different network sizes

Application	P_{laser}^{tot} (mW)				Run time			
	4 × 4	6 × 6	8 × 8	10 × 10	4 × 4	6 × 6	8 × 8	10 × 10
FFT-1024_complex	1096	6333	13549	28040	10.05 s	41.01 s	2.63 min	3.81 min
Fpppp	1064	2756	3617	4243	0.55 s	1.16 min	2.70 min	9.32 min
H264-1080p_dec	468	2290	3850	3531	24.12 s	26.33 s	29.07 s	6.08 min
Sparse	200	251	204	206	42.98 s	1.67 min	3.44 min	15.11 min

Table 8. Characteristics of the applications analyzed

Application	CG completeness			
	4 × 4	8 × 8	12 × 12	16 × 16
FFT-1024_complex	0.9292	0.6173	0.3150	0.1158
FPPPP	0.8917	0.1664	0.0382	0.0117
H264-1080p_dec	0.3958	0.1766	0.0367	0.0284
SPARSE	0.175	0.0094	0.0020	0.0005

consumption through an application-specific mapping optimization. We showed that the laser power consumption is highly dependent on the mapping choices. Then, we proposed some design space exploration algorithms which automatically map the application cores to the NoC tiles while minimizing the laser power consumption. Experimental studies show that the laser power consumption can be significantly reduced allowing an enhanced energy efficiency.

ACKNOWLEDGMENTS

The work has been partially supported by the European Commission under Grant No.: 671668.

REFERENCES

[1] Johnnie Chan, Gilbert Hendry, Keren Bergman, and Luca P Carloni. 2011. Physical-layer modeling and system-level design of chip-scale photonic interconnection networks. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on* 30, 10 (2011), 1507–1520.

[2] Chao Chen, Tiansheng Zhang, Pietro Contu, Jonathan Klamkin, Ayse K Coskun, and Ajay Joshi. 2014. Sharing and placement of on-chip laser sources in silicon-photonic NoCs. In *Networks-on-Chip (NoCS), 2014 Eighth IEEE/ACM International Symposium on.* IEEE, 88–95.

[3] Alessandro Cilardo and Edoardo Fusella. 2016. Design automation for application-specific on-chip interconnects: A survey. *Integration, the VLSI Journal* 52 (2016), 102–121.

[4] Luan HK Duong, Zhehui Wang, Mahdi Nikdast, Jiang Xu, Peng Yang, Zhifei Wang, Zhe Wang, Rafael KV Maeda, Haoran Li, Xuan Wang, et al. 2016. Coherent and incoherent crosstalk noise analyses in interchip/intrachip optical interconnection networks. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 24, 7 (2016), 2475–2487.

[5] Edoardo Fusella and Alessandro Cilardo. 2016. Crosstalk-aware automated mapping for optical networks-on-chip. *ACM Transactions on Embedded Computing Systems (TECS)* 16, 1 (2016), 16.

[6] Edoardo Fusella and Alessandro Cilardo. 2016. Lighting Up On-Chip Communications with Photonics: Design Tradeoffs for Optical NoC Architectures. *Circuits and Systems Magazine, IEEE* 16, 3 (thirdquarter 2016), 4–14.

[7] Edoardo Fusella and Alessandro Cilardo. 2016. Minimizing power loss in optical networks-on-chip through application-specific mapping. *Microprocessors and Microsystems* 43 (2016), 4–13.

[8] Edoardo Fusella and Alessandro Cilardo. 2016. PhoNoCMap: an application mapping tool for photonic networks-on-chip. In *2016 Design, Automation & Test in Europe Conference & Exhibition (DATE).* IEEE, 289–292.

[9] Edoardo Fusella and Alessandro Cilardo. 2017. H^2ONoC: A Hybrid Optical-Electronic NoC Based on Hybrid Topology. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 25, 1 (Jan 2017), 330–343. https://doi.org/10.1109/TVLSI.2016.2581486

[10] Michael R Garey and David S Johnson. 2002. *Computers and intractability.* Vol. 29. wh freeman New York.

[11] Martijn JR Heck and John E Bowers. 2014. Energy efficient and energy proportional optical interconnects for multi-core processors: Driving the need for on-chip sources. *IEEE Journal of Selected Topics in Quantum Electronics* 20, 4 (2014), 332–343.

[12] Jingcao Hu and Radu Marculescu. 2005. Energy-and performance-aware mapping for regular NoC architectures. *Computer-Aided Design of Integrated Circuits and Systems, IEEE Transactions on* 24, 4 (2005), 551–562.

[13] George Kurian, Chen Sun, Chia-Hsin Owen Chen, Jason E Miller, Jurgen Michel, Lan Wei, Dimitri A Antoniadis, Li-Shiuan Peh, Lionel Kimerling, Vladimir Stojanovic, et al. 2012. Cross-layer energy and performance evaluation of a nanophotonic manycore processor system using real application workloads. In *Parallel & Distributed Processing Symposium (IPDPS), 2012 IEEE 26th International*. IEEE, 1117–1130.

[14] Sébastien Le Beux, Hui Li, Ian O'Connor, Kazem Cheshmi, Xuchen Liu, Jelena Trajkovic, and Gabriela Nicolescu. 2014. Chameleon: Channel efficient optical network-on-chip. In *Design, Automation and Test in Europe Conference and Exhibition (DATE), 2014*. IEEE, 1–6.

[15] Benjamin G Lee, Aleksandr Biberman, Po Dong, Michal Lipson, and Keren Bergman. 2008. All-optical comb switch for multiwavelength message routing in silicon photonic networks. *Photonics Technology Letters, IEEE* 20, 10 (2008), 767–769.

[16] Weichen Liu, Jiang Xu, Xiaowen Wu, Yaoyao Ye, Xuan Wang, Wei Zhang, Mahdi Nikdast, and Zhehui Wang. 2011. A NoC Traffic Suite Based on Real Applications. In *VLSI (ISVLSI), 2011 IEEE Computer Society Annual Symposium on*. IEEE, 66–71.

[17] Gianlorenzo Masini, Giovanni Capellini, Jeremy Witzens, and Cary Gunn. 2007. A four-channel, 10 Gbps monolithic optical receiver in 130nm CMOS with integrated Ge waveguide photodetectors. In *National Fiber Optic Engineers Conference*. Optical Society of America, PDP31.

[18] Srinivasan Murali and Giovanni De Micheli. 2004. Bandwidth-constrained mapping of cores onto NoC architectures. In *Proceedings of the conference on Design, automation and test in Europe-Volume 2*. IEEE Computer Society, 20896.

[19] Gabriela Nicolescu, Sebastien Le Beux, Mahdi Nikdast, and Jiang Xu. 2017. Photonic Interconnects for Computing Systems. (2017).

[20] Mahdi Nikdast, Jiang Xu, Luan Huu Kinh Duong, Xiaowen Wu, Xuan Wang, Zhehui Wang, Zhe Wang, Peng Yang, Yaoyao Ye, and Qinfen Hao. 2015. Crosstalk noise in WDM-based optical networks-on-chip: A formal study and comparison. *IEEE Transactions on Very Large Scale Integration (VLSI) Systems* 23, 11 (2015), 2552–2565.

[21] IM Oliver, DJd Smith, and John RC Holland. 1987. Study of permutation crossover operators on the traveling salesman problem. In *Genetic algorithms and their applications: proceedings of the second International Conference on Genetic Algorithms: July 28-31, 1987 at the Massachusetts Institute of Technology, Cambridge, MA*. Hillsdale, NJ: L. Erlhaum Associates, 1987.

[22] Yan Pan, John Kim, and Gokhan Memik. 2010. Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar. In *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*. IEEE, 1–12.

[23] Cédric Robert, T Nguyen Thanh, A Létoublon, M Perrin, C Cornet, C Levallois, JM Jancu, J Even, P Turban, A Balocchi, et al. 2013. Structural and optical properties of AlGaP confinement layers and InGaAs quantum dot light emitters onto GaP substrate: Towards photonics on silicon applications. *Thin Solid Films* 541 (2013), 87–91.

[24] Günther Roelkens, Liu Liu, Di Liang, Richard Jones, Alexander Fang, Brian Koch, and John Bowers. 2010. III–V/silicon photonics for on-chip and intra-chip optical interconnects. *Laser & Photonics Reviews* 4, 6 (2010), 751–779.

[25] Pradip Kumar Sahu and Santanu Chattopadhyay. 2013. A survey on application mapping strategies for network-on-chip design. *Journal of Systems Architecture* 59, 1 (2013), 60–76.

[26] Assaf Shacham, Keren Bergman, and Luca P Carloni. 2008. Photonic networks-on-chip for future generations of chip multiprocessors. *Computers, IEEE Transactions on* 57, 9 (2008), 1246–1260.

[27] Eric Tournié, Laurent Cerutti, Jean-Baptiste Rodriguez, Huiyun Liu, Jiang Wu, and Siming Chen. 2016. Metamorphic III–V semiconductor lasers grown on silicon. *MRS Bulletin* 41, 03 (2016), 218–223.

[28] Xiaowen Wu, Jiang Xu, Yaoyao Ye, Zhehui Wang, Mahdi Nikdast, and Xuan Wang. 2014. SUOR: Sectioned undirectional optical ring for chip multiprocessor. *ACM Journal on Emerging Technologies in Computing Systems (JETC)* 10, 4 (2014), 29.

[29] Yiyuan Xie, Mahdi Nikdast, Jiang Xu, Wei Zhang, Qi Li, Xiaowen Wu, Yaoyao Ye, Xuan Wang, and Weichen Liu. 2010. Crosstalk noise and bit error rate analysis for optical network-on-chip. In *Proceedings of the 47th Design Automation Conference*. ACM, 657–660.

[30] Yaoyao Ye, Jiang Xu, Xiaowen Wu, Wei Zhang, Weichen Liu, and Mahdi Nikdast. 2012. A torus-based hierarchical optical-electronic network-on-chip for multiprocessor system-on-chip. *ACM Journal on Emerging Technologies in Computing Systems (JETC)* 8, 1 (2012), 5.