

Document downloaded from:

<http://hdl.handle.net/10251/72599>

This paper must be cited as:

Puche Lara, J.; Lechago Buendia, S.; Petit Martí, SV.; Gómez Requena, ME.; Sahuquillo Borrás, J. (2016). Accurately Modeling a Photonic NoC in a Detailed CMP Simulation Framework. IEEE. doi:10.1109/HPCSim.2016.7568361.



The final publication is available at

<http://dx.doi.org/10.1109/HPCSim.2016.7568361>

Copyright IEEE

Additional Information

© 2016 IEEE. Personal use of this material is permitted. Permission from IEEE must be obtained for all other uses, in any current or future media, including reprinting/republishing this material for advertising or promotional purposes, creating new collective works, for resale or redistribution to servers or lists, or reuse of any copyrighted component of this work in other works.

Accurately Modeling a Photonic NoC in a Detailed CMP Simulation Framework

José Puche*, Sergio Lechago†, Salvador Petit*, María E. Gómez* and Julio Sahuquillo*

*Departamento de Ingeniería de Sistemas y Computadores

Universidad Politécnica de Valencia

Email: jopucla@posgrado.upv.es

†Centro de Tecnología Nanofotónica de Valencia

Universidad Politécnica de Valencia

Email: serlecbu@ntc.upv.es

Abstract—Photonic interconnects are a promising solution for the so-called communication bottleneck in current Chip Multiprocessor (CMPs) architectures. This technology presents an inherent low-latency and power consumption almost independent of communication distance, which are really desirable features in future Networks on Chip for next CMPs generations. However, since nanophotonic technology is still growing and therefore in an immature state, current simulators of detailed systems may not provide accurate models of photonic components. In this context, non-representative results are obtained when unaccurate photonic models are assumed.

This paper summarizes all of the components that conform a fully operative photonic NoC and presents their current state of the art. Moreover, we evaluate a realistic photonic network that consists of two photonic rings and a token-based arbitration mechanism and compare it against a non-realistic model. In addition, both realistic and non-realistic schemes are evaluated under different configurations varying the number of wavelengths that photonic waveguides employ. The experimental results show that the non-realistic NoC presents up to $6\times$ network latency deviation with respect to the accurate model. This deviation is translated into a performance deviation higher than 10% in several applications studied, which demonstrates the importance of accurate models when simulating current technologies under development like nanophotonics.

Finally, a power consumption model of the realistic photonic network is presented. The results show that the overall photonic network power consumption grows with the number of wavelengths per waveguide since the number of required modulators and receivers becomes higher. In this way, the proposed realistic photonic network, which employs only two wavelengths for arbitration and destination selection tasks, increases its power consumption up to 3%, so network designs with more complex arbitration mechanisms must take into account the impact of the number of wavelengths on the power consumption.

I. INTRODUCTION

In order to keep pace with Moore’s law, microprocessors have leveraged multicore architectures throughout the last ten years. Technology advances currently enable integrating hundreds of cores on a single chip, so rising the potential computational power [?]. These Chip Multiprocessors (CMPs), however, require from efficient on-chip interconnection networks (NoC) to provide fast communication among cores, caches, and the memory controllers. Otherwise, the NoC could seriously limit the potential performance associated to these architectures. Moreover, the increased number of cores and its

communication costs must be designed to be within a limited power budget.

Future CMP generations need to face the challenges of efficient global communications and low power consumption. In this context, current NoCs are not likely to properly face these design issues for future CMPs because of electrical technology constraints, especially when the number of cores increases and some messages must traverse long distances (i.e., due to a high number of hops) to reach their destinations.

Electrical networks used in CMP architectures often present a mesh topology where memory controllers (MC) are usually placed at the corners and edges of the processor chip. Since the NoC must interconnect all the tiles, the higher the number of cores, the higher the average distance from the nodes to main memory. In such a scheme, when a node requests a cache block, associated latencies and energy consumption can be unacceptable depending on the distance to the MC. Current multicore architectures incorporate several MCs to alleviate this drawback but, unfortunately, the number of MCs can not scale linearly with the number of nodes [1].

The need of a faster on-chip multicore communication technology has led to lay out the use of CMOS-compatible photonic interconnects as a possible solution. Nanophotonics technology has experienced a vertiginous development during the last decade and this trend is expected to continue in the next years. Because of its high bandwidth and its low energy cost, which scarcely varies with the communication distance, CMOS compatible photonics interconnects is the most promising technology to satisfy future CMPs’ bandwidth demands [2].

Manycore and multicore architectures can leverage the capacities provided by nanophotonics to reduce their network latency and, as a result, their memory access cycles. However, nanophotonic technology is still maturing, which difficults an up-to-date modeling in CMP simulation frameworks, mainly developed by computer architects whose focus is on computing and communications aspects. This situation leads to inaccurate models whose results could present important deviations. Therefore, detailed and accurate simulation environments are needed to develop reliable CMP simulators considering nanophotonic interconnects.

Nanophotonics technology presents a variety of brand new components that must be modeled and incorporated into the existing CMP simulation frameworks. In this paper we present a photonic network model that introduces all of the components a real photonic network employs. We explain how each one of these components affects the network performance and how the whole system power consumption can be compromised. Also we show how different simulation configurations alter the execution time of the studied workloads to demonstrate that an unrealistic modeling produces unreliable results.

The remainder of this paper is organized as follows. Section II presents some nanophotonics background, focusing on the behavior of its components. Section III provides an overview of the current photonic technology state of the art. Section IV introduces the modeling of the photonic components. Section V describes the environment where experimental evaluations take place. Section VI discusses the experimental evaluation results. Finally, Section VII presents some concluding remarks.

II. BACKGROUND

Advances in silicon nanophotonics currently allow the development of a complete functional optical network in a single chip [?]. The optical network operation requires several photonic components, whose functions are briefly described in this section.

A laser source is required to inject light into the chip. This light is carried by *waveguides* to the rest of the components of the optical network. In addition, the laser light is multiplexed into different *wavelengths*. To do so, the *Dense Wavelength Division Multiplexing* (DWDM) technique is used. DWDM allows several network nodes to communicate at the same time through multiple wavelengths in a single waveguide [3].

The DWDM technique requires resonators to separate the different wavelengths that compose the light signal. A resonator is a little ring that filters a given wavelength from the waveguide. By default, the wavelength filtered by a resonator is determined by its ring diameter, which usually ranges from 3 to 5 μm [4]. A resonator can be tuned to filter a different wavelength by applying an electric charge to it or by rising its temperature. Note that resonators need to be tuned every time that they are used to establish a transmission in a different wavelength than their default one. Thus, the network power consumption depends on the number of resonators and transmissions.

Resonators are used both in the source and the destination nodes of a transmission. In the source node, an electro-optical modulator conveys the digital signal to be transmitted inside the wavelength filtered by the resonator. In the destination node, a resonator tuned to the same wavelength filters the optical signal and guides it to a receiver or photodetector, which finally transforms the optical signal into an electrical signal that can be used in a digital circuit. Notice that both electrical-to-optical and optical-to-electrical conversion times must be taken into account when evaluating the performance of a photonic network.




Fig. 1. End-to-end transmission between two network nodes using photonic interconnects.

Figure 1 shows an example of an end-to-end photonic transmission between two network nodes A and B. To transmit a bit flow from node A to node B, first, the resonator on node A absorbs a given λ_i wavelength from the injected laser light. Then the bit flow is modulated in the filtered optical signal, which in turn is brought to the waveguide, where it is routed to node B. On node B, the resonator filters the same wavelength, allowing the B receiver to react to the corresponding optical signal and convert it to a bit flow to be processed in node B.

III. STATE OF THE ART IN PHOTONIC TECHNOLOGY

Current research efforts focused on Photonic Integrated Circuits (PICs) are concentrated on the realization of reliable hybrid silicon lasers, electro-optical modulators and receivers, the most critical building blocks of photonic circuits. The promising research results pave the way to fully on-chip integrated devices able to overcome inherent limitations of electronic performances.

Laser sources are the most difficult devices to be integrated on-chip due to power, area, and optical signal attenuation constraints. Duan et al. have developed hybrid silicon/III-V lasers exhibiting new features and lower power consumption than previous works [5], [6]. However, these advances do not yet achieve the ultra-low power consumption required for on-chip laser integration. Moreover, integrated lasers are only able to provide output signal powers of tens of mW, raising attenuation concerns. Nevertheless, it is expected that in the next few years higher figures will be accomplished, allowing the exploitation of on-chip lasers real potential. In fact, some works on photonic NoCs assume that on-chip lasers will be integrated in future technologies since they are much more energy-efficient [18].

The switching capacity of electro-optical modulators represents the key feature for establishing the operation bandwidth of any PIC. High-bandwidth modulation in silicon (achieving up to 3050 Gb/s data rates and working at switching times of several GHz [8], [9]) can be realized by free-carrier induced index change and using biased p-n structures (carrier depletion) [7].

Optical coherent receivers, which convert the amplitude, phase, and polarization of an optical field signal into the electrical domain have already been integrated showing similar performance that those yielded by commercial devices with

very high data conversion (up to 224 Gb/s with PDM-16-QAM signals) [10].

Regarding current research on other PIC components, critical issues are to minimize light signal attenuation in the manufacturing process of waveguides [?] and to reduce the width of the light spectrum that resonators can filter. The latter characteristic defines the number of wavelengths achievable by DWDM, which currently ranges from 64 to 160 wavelengths per waveguide.

With respect to state-of-the-art research on hybrid photonic-electronic and pure photonic NoC implementations, Vantrease et. al explore future manycores bandwidth requirements and propose Corona [12], a manycore 3D architecture that employs photonic technology for both on-chip and off-chip communication. Kurian et. al present ATAC [13], a 1K-core system that communicates through a photonic NoC. In [14], authors propose Firefly, a hybrid NoC that leverages node clustering and employs photonic interconnects for inter-cluster communications. Finally, FlexiShare [15] uses a photonic ring combined with DWDM to interconnect a 64-core CMP.

Research on photonic NoCs is closely related to research on DWDM arbitration techniques. Since most DWDM-based communication schemes require wavelength sharing, some works consider the arbitration as an important part of the communication efficiency [12], [15]. In this context, Vantrease et. al [16] also leverage photonic technology to perform arbitration related tasks. They identify latency, average network utilization and fairness as the key features a suitable arbitration mechanism must address.

IV. MODELING PHOTONIC COMPONENTS

Every component mentioned in Section II is properly modeled in our simulation environment. For each component, we identify the critical properties that affect system performance and energy consumption.

Regarding the laser, we assume that is placed off-chip due to the issues explained Section III. The power budget of an off-chip laser can vary from 1W to even more than 5W. The exact laser wattage depends on waveguide characteristics such as refractive indexes, turns, couples, splits, etc.

With respect to waveguides, they have two main properties that impact the communication latency: i) the refractive index and ii) the optical path length. Waveguides usually employed in prototypes are made of crystalline silicon and silicon oxide, which present a refractive index of 3.4401 and 1.4298, respectively. This leads to an average index of 2.439. As a result, the propagation of light speed over the silicon die is assumed to be $12.3mm/100ps$. Regarding the optical path length, it depends on the number of interconnected nodes and the chip dimensions. In our baseline system (see Section V), a $576 mm^2$ CMP [16] formed by 16 tiles and one memory controller (*i.e.* a network that interconnects up to 17 nodes) requires a 116 mm path length.

Modulators and receivers also impact the communication latency. In our simulation environment, both modulators and receivers are modeled as components that take 1 cycle at

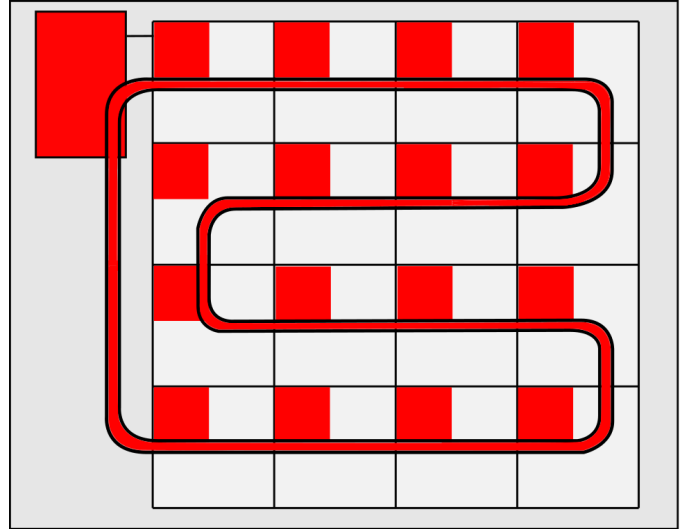


Fig. 2. Schematic of a 116 mm length waveguide on a 16-core tiled CMP.

a given frequency to transmit 1 bit. Since state-of-the-art modulators present a switching time by 100 picoseconds, we set the modulators frequency to 10 GHz. Regarding receivers, we found in the literature that their latency ranges from ten to hundreds of picoseconds [18]. Thus, we assume that the latency of receivers matches the latency of modulators.

The DWDM technology allows different communication schemes to leverage multiple wavelengths as a shared transmission medium. In general, there are four well-known DWDM communication schemes [3]: Single Writer Single Reader (SWSR), Single Writer Multiple Reader (SWMR), Multiple Writer Single Reader (MWSR) and Multiple Writer Multiple Reader (MWMR). Apart from the simple SWSR scheme, all the schemes require arbitration to access the wavelengths [16]. Therefore, in order to correctly model a photonic NoC it is necessary to define both the communication scheme and the arbitration mechanism. In this context, token-based arbitration approaches are typically used when evaluating and modeling communication schemes, and token injection and extraction latencies are also properly considered in our simulation environment.

V. SYSTEM OVERVIEW

Figure 2 shows a block diagram of the baseline system: a 16 tiled CMP where all nodes and the memory controller are connected by a single photonic ring. Each tile consists of an out-of-order core with L1 and L2 private caches and the network interface to access the ring.

The photonic ring is composed of 2 waveguides. One of the waveguides is devoted to messages sent from the cores to the memory controller while the other waveguide is used for transmissions in the opposite direction. We use the MWSR communication scheme when sending requests to the memory controller, which implies that two nodes cannot perform a request to the memory controller at the same time. Therefore,

Processing Core	
Number of cores	16
Frecuency	3GHz
Issuing policy	Out of order
Branch predictor	bimodal/gshare hybrid: gshare with 14-bits global history + 16K 2-bit counters, bimodal with 4K 2-bit counters, and selection with 4K 2-bit counters
Issue/Commit width	4 instructions/cycle
ROB size	256 entries
Memory hierarchy	
L1 Instruccion cache	Private, 32KB, 8 ways, 64Bytes-line, 2 cycles
L1 Data cache	Private, 32KB, 8 ways, 64Bytes-line, 2 cycles
L2	Private, 256KB, 16 ways, 64Bytes-line, 11 cycles
Photonic Ring	
Topology	Ring
Waveguides	2
Wavelengths	64 wavelengths per waveguide
Frequency	10 GHz
Modulator lat	1 cycle
Photodetector lat	100 ps
Arbitration	Token channel
Phit size	64 bits
Roundtrip lat	14 cycles on idle

TABLE I
CONFIGURATION OF THE SIMULATED SYSTEM.

optical arbitration between multiple cores is required. Optical arbitration is driven via a wavelength-routed token scheme, which consists of passing a single token among all possible senders (i.e., the cores) [16]. Notice that even when only one core is transmitting to the memory controller, the token must be released regularly to check if other there are new cores ready to transmit.

On the other hand, messages from the memory controller are sent following a SWMR communication scheme. This scheme avoids arbitration delay but it requires the tuning of the destination resonators before actually sending the message. To this end, one wavelength is used to activate the destination core before transmission.

Table I summarizes the main baseline system parameters. Photonic ring latency scales with the length of the optical path. Assuming a squared N -core CMP die, the area occupied by each core has a length and width proportional to $1/\sqrt{N}$. Therefore, a 116 mm waveguide length is required to allow each node to reach the memory controller in the 16-core 576 mm^2 die processor.

Once arbitration has been performed, the overall ring roundtrip latency accounts for modulator, waveguide and receiver latencies. The waveguide latency is defined by the optical path length and the light propagation speed over silicon, as explained in Section IV. Thus, the 10GHz photonic ring roundtrip is 12 cycles plus conversion times. As pointed out in Section IV, these conversion times consists of 1 modulation cycle and 1 reception cycle. As a result, the roundtrip latency becomes 14 cycles.

The number of wavelengths in which laser light can be multiplexed mainly depends on the photonic technology. Previous works on photonic networks assume 64 wavelengths per waveguide [12], [14], [15], which is the value typically used.

However, some recent works point out that this figure can grow over 160 wavelengths per waveguide in state-of-the-art technologies [18]. To make the study up-to-date, this paper evaluates the benefits of photonics considering both 64 and 160 wavelengths per waveguide.

VI. EXPERIMENTAL RESULTS

Multiple photonic approaches have been modeled and evaluated with the Multi2Sim simulation framework [19], which simulates in detailed the out-of-order cores and the memory hierarchy. The Multi2Sim network layer has been widely extended to properly model the photonic NoC and every nanophotonic component described in previous sections.

Experiments have been performed using the SPEC2006 benchmark suite. We show executions of both individual applications and multiprogrammed mixes to explore how the detailed network simulation affects the obtained results. Each application in the mix commits at least 100M instructions after fastforwarding the initial 300M instructions. This fastforward is done to warm up caches and to avoid performance differences due to this reason. Moreover, when evaluating the performance of multiprogrammed mixes, all the applications keep running until the last benchmark finishes the target number of instructions. Otherwise, the fastest benchmarks would be more affected by contention than the slowest ones.

A. Benchmark characterization

As a first step we analyze and characterize the memory activity of each individual benchmark. Benchmarks which present a high number of misses in their last level cache (the L2 cache in our system) access the interconnection network and memory more frequently. The system employs 8-byte packets for requests and acknowledgments and 72-byte packets for data messages, which are composed by an 8-byte header and a 64-byte data payload.

Figure 3 shows the number of Memory Accesses per Kilo-Instructions (MAPKI) performed by the studied benchmarks in increasing value order. At first glance, we can distinguish two different groups of applications. Applications on the left side present a really low number of memory accesses (e.g., MAPKI=2), hence being their performance is not significantly affected by the NoC. These applications can store almost their




Fig. 3. Memory Accesses Per KiloInstructions (MAPKI) of SPEC2006.




Fig. 4. Network latency of benchmarks executed with and without arbitration delay and 64 wavelengths per waveguide.

entire working set in the private caches and thus they scarcely access to main memory. On the other hand, applications on the right side incur on a high number of memory accesses.

Since we are interested in showing how photonic network modeling affects experimental results, from now on we will differentiate these two kinds of applications when discussing benchmarks behavior.

B. Individual execution

Different scenarios have been considered to evaluate how detailed photonic simulation impacts on the achieved performance. To clearly differentiate the impact of the photonic network on individual application performance, we first execute every application in an isolated way. Nevertheless, since arbitration is a key characteristic of photonic networks, we take arbitration overheads into account even if the executed application is alone in the system. This means that in the experiments performed in this section, before sending a message to the memory controller, the executing core must wait for a full ring token roundtrip. This emulates the overhead of checking if there are more cores in the ring ready to send a message and releasing the token to avoid their starvation. Second, the impact on performance of the number of wavelengths per waveguide supported by current photonic state-of-the-art is also studied. This number affects the photonic ring bandwidth so we explore how a higher number of wavelengths per waveguide can attenuate network delays.

Figure 4 shows the average network latency of the studied applications with and without arbitration delay and assuming 64 wavelengths per waveguide. The lower frame of the bars refers to the NoArb64 scheme which does not model arbitration delays, while the upper frame (Arb64) represents the latency added by arbitration. In the NoArb64 scheme, messages are not delayed by the mentioned token roundtrip latency. As it can be seen, network latency for this configuration is quite homogeneous and close to 5 network cycles for all the evaluated benchmarks. This value grows over 35 cycles when arbitration delay is modeled, which means that modeling arbitration latency introduces a network latency deviation higher than a 6× factor on average.




Fig. 5. Network latency of benchmarks executed with and without arbitration delay and 160 wavelengths per waveguide.

Figure 5 shows results corresponding to the same arbitration schemes but with a NoC including 160 wavelengths per waveguide. This helps the network to reduce the cycles needed to send 72-byte messages through the ring but the latency corresponding to arbitration remains constant. The results obtained in this configuration show that arbitration delay represents a high percentage of the average network latency so increasing the available bandwidth can not significantly reduce the overall latency.

Results presented in Figure 4 and Figure 5 point out the potential latency deviation that experiments can incur with non-realistic models. Applications as *leslie3d*, *astar* or *bzip2* present an average network latency variation as high as a 10× factor between executions with and without arbitration delays.

Figure 6 shows the performance results for the different configurations. Applications are shown by increasingly order of MAPKI, hence the IPC of the benchmarks on the right side of the plot is lower than that of the applications on the left side, since the former applications perform a higher use of the network.

Performance deviations incurred by not modeling arbitration delays in both 64 and 160 wavelengths configurations is shown in Figure 7, which plots the IPC increase experienced when arbitration overheads are not taken into account. As it can be seen, non-realistic arbitration modeling significantly affects the obtained IPC results. As expected, applications on the right side show pronounced IPC variations between the two arbitration schemes while applications on the left side remain with a similar performance. This figure outlines that an unrealistic arbitration model can suppose an error in system performance results as large as 14%, depending on the number of memory accesses the application performs. In contrast, varying the number of wavelengths per waveguide does not incur on relevant performance growth.

C. Multiprogrammed workloads execution

Results shown in Section VI-B demonstrate the deviation suffered by the applications because of the lack of accuracy in the photonic model in isolated execution. In contrast, this




Fig. 6. Absolute IPC of executions with and without arbitration delay in 64 wl/wg and 160 wl/wg configurations.




Fig. 7. IPC deviation of executions with and without arbitration delay in 64 wl/wg and 160 wl/wg configurations.

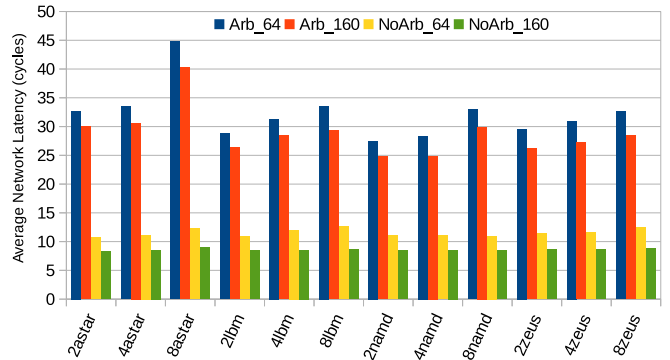


Fig. 8. Network latencies of executions with and without arbitration delay in 64 wl/wg and 160 wl/wg configurations.

section presents results corresponding to applications executed concurrently with co-runners in other cores of the CMP, so NoC and memory contention are introduced in these experiments. Aiming to avoid introducing deviations due to different co-runners, these are just several instances of the studied workload. Thus, configurations as 2astar, 4astar, etc. study network and system performance when 2 and 4 instances of astar are co-running. In this way, the impact of contention on average network latency and system performance is studied as it can hide the deviations shown in the results obtained during individual execution.

Figure 8 summarizes network latencies for both 64 and 160 wavelengths per waveguide configurations. Results obtained present a similar behaviour as the obtained during individual execution. Because of the contention effect, average network latency increases with the number of corunners. However, this effect does not hide arbitration-related deviations in network latency since schemes with and without arbitration still vary from a 3× to a 4.5× factor. These deviations also occur with 160 wavelengths per waveguide.

Figure 9 shows the deviation in the system performance that rises when arbitration is not considered. IPC growth due to lack of arbitration in these executions highly depends on the application. In the case of namd, the deviation is almost null since this application does not perform a significant number

of memory accesses. Regarding astar and lbm, both applications present similar behaviour. As the network contention grows, the IPC growth experimented by the workloads in non-arbitration approaches increases.

Finally, results associated to zeusmp show that this application is more affected by memory contention than by network contention. When this application is executed in isolation or with only one corunner its IPC remains almost constant regardless of the studied arbitration approach. However, as contention grows, its IPC falls. Therefore, any reduction caused by not correctly modeling the arbitration overheads results in relative IPC increase (see results of 4zeus and 8zeus). As a result, applications as zeusmp can hide these network latency variations when executed in isolation, so contention must be considered to get accurate and representative results.

D. Power consumption

This section presents the power consumption associated to the photonic network model discussed in Section V. In order to achieve a detailed power model, every photonic component is studied separately and its power consumption based on current state of the art studies and predictions is obtained. The energy consumption model employed distinguishes between dynamic




Fig. 9. IPC deviation of executions with and without arbitration delay in 64 wl/wg and 160 wl/wg configurations.

TABLE II
ENERGY CONSUMPTION PARAMETERS.

Energy_dynamic	
Transmitters	135 fJ/bit
Receivers	365 fJ/bit
Total	0.42 pJ/bt
Energy_static	
Laser power output	22.5 mW
Microrings Tuning	1.35 mW/ring
Total	22.5 + (1.35 x rings) mW

energy consumption caused by modulators (transmitters) and photodetectors (receivers) and static energy consumption due to tuning of microrings associated to these components and laser consumption. In Section IV we assumed an off-chip laser taking into account the existent proposals. In contrast, if an on-chip laser is used, then the power model should also take into account its static consumption. To perform a self-contained analysis, this section includes on-chip laser energy budget.

Regarding dynamic energy consumption, recent studies have achieved efficient transmitters and receivers developments [?]. Their power consumption expressed by fJ per bit is shown in Table II. Current research on photonics is focused on the power consumption associated to different NoC components. However, an extended consumption model should also account the energy consumed by the remaining electronic logic needed to interface with the waveguide.

Static energy consumption is defined by the laser and microrings tuning. Notice that the energy consumption accounted due to this tuning depends on the number of microrings the photonic network employs. Thus, the higher number of rings the higher energy is required. The number of photonic microrings is closely related to the selected communication scheme (MWSR and SWMR in this case) and the number of wavelengths used. MWSR and SWMR schemes imply that for a given number x of wavelengths, every node must include $2x$ microrings (x to send and x to receive). In our model, this means that every node should have 64 microrings on each ring, as well as receptors should present the same number in order to receive light in the corresponding wavelengths.




Fig. 10. Static power consumption (mW) required by the four photonic network configurations studied.

The energy consumption relative to thermo-optic microring tuning depends on the number of channels that the microring is able to filter and the channel spacing as well. Some works as [?] point out that the effective tuning efficiency is calculated to be $27\mu W/GHz/ring$. We assume 50 GHz of channel spacing this means that power consumption relative to microring tuning is 1.35 mW/ring, as shown in Table II.

Laser power consumption depends mostly on the losses that the optical path can introduce into the light signal. Moreover, different kinds of lasers with a wide offer in terms of power and energy consumption are available in the current state of the art. In this section, we assume an on-chip hybrid silicon laser which presents an injection power of 22.5 mW [18].

Figure 10 shows the total amount of static power consumption required by the studied schemes. Photonic rings that employ 160 wavelengths increase their power consumption due to the higher number of microrings. Future configurations with higher number of nodes and wavelengths should face the challenge of reducing the power consumption associated to thermal tuning while keeping bandwidth rates below acceptable levels.

Due to the use of MWSR, the arbitration technique only requires one wavelength to pass the token between senders because there is just one possible destination node for them. In this context, only one resonator must be added to each node to guarantee that the token is properly transmitted. In addition, destination selection performed for the messages sent by memory controllers in SWMR requires one wavelength as well. These additional wavelengths do not seriously affect the overall power consumption, since they only suppose about 3% and 1.23% of energy increase with 64 and 160 wavelengths respectively. However, other communication schemes as MWMR or configurations with multiple destination nodes that require higher number of tokens should take power consumption associated to arbitration techniques into account. Otherwise, power consumption could experiment an exponential growth as the number of wavelengths required for arbitration is increased.

VII. CONCLUSIONS

An analysis about the importance of employing accurate models when simulating novel technologies as nanophotonics is performed in this paper. We identify how every component that conforms a fully operative photonic NoC should be modeled in order to obtain accurate and representative results both in network and system performance. In this paper we also have reviewed current state of the art in photonic technology and pointed out realistic and future parameters that some photonic components as waveguides, modulators or photodetectors present.

Aiming to quantify the deviation that a wrong photonic model could present in a detailed simulation environment, we have modeled and evaluated a realistic proposal based in two photonic rings and a simple arbitration technique and compare it against a non-realistic configuration with no arbitration. Experimental results outline that the variation in average network latency between the two approaches can be as high as 1000%, which is translated in an IPC deviation higher than a 10% in some cases. Moreover, according to current state of the art power consumption model, the realistic approach increases the overall network energy consumption up to 3% respect to the non-realistic setup, and this deviation can grow as the complexity of the arbitration technique or the number of wavelengths increases. These deviations demonstrate that current simulators must be properly extended with reliable and accurate models in order to obtain representative results when researching on novel and immature technologies.

ACKNOWLEDGMENTS

This work was supported by the Spanish Ministerio de Economía y Competitividad (MINECO) and by Plan E funds under Grant TIN2015-66972-C5-1-R and the ExaNest project, funded by the European Union's Horizon 2020 research and innovation programme under grant agreement No 671553.

REFERENCES

- [1] S. R. Vangal, J. Howard, G. Ruhl, S. Dige, H. Wilson, J. Tschanz, D. Finan, A. Singh, T. Jacob, N. Borkar, and S. Borkar, "An 80-tile sub-100-w teraflops processor in 65-nm cmos," Oct 2010.
- [2] D. Miller, "Device requirements for optical interconnects to silicon chips," *Proceedings of the IEEE*, vol. 97, no. 7, pp. 1166–1185, July 2009.
- [3] A. C. Bergman K. Carloni, L. P. Bibermani and H. G., *Photonic Network-on-Chip Design*. Springer-Verlag New York, 2014, vol. 68, no. 1.
- [4] Q. Xu, D. Fattal, and R. G. Beausoleil, "Silicon microring resonators with 1.5- μm radius," *Opt. Express*, vol. 16, no. 6, pp. 4309–4315, Mar 2008. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-16-6-4309>
- [5] G.-H. Duan, J.-M. Fedeli, S. Keyvaninia, and D. Thomson, "10 gb/s integrated tunable hybrid iii-v/si laser and silicon mach-zehnder modulator," in *European Conference and Exhibition on Optical Communication*. Optical Society of America, 2012, p. Tu.4.E.2. [Online]. Available: <http://www.osapublishing.org/abstract.cfm?URI=ECEOC-2012-Tu.4.E.2>
- [6] G. H. Duan, C. Jany, A. L. Liepvre, M. Lamponi, A. Accard, F. Poingt, D. Make, F. Lelarge, S. Messaoudene, D. Bordel, J. M. Fedeli, S. Keyvaninia, G. Roelkens, D. V. Thourhout, D. J. Thomson, F. Y. Gardes, and G. T. Reed, "Integrated hybrid iii x2013;v/si laser and transmitter," in *Indium Phosphide and Related Materials (IPRM), 2012 International Conference on*, Aug 2012, pp. 16–19.

- [7] R. Soref and B. Bennett, "Electrooptical effects in silicon," *IEEE Journal of Quantum Electronics*, vol. 23, no. 1, pp. 123–129, Jan 1987.
- [8] A. Liu, L. Liao, D. Rubin, H. Nguyen, B. Ciftcioglu, Y. Chetrit, N. Izhaky, and M. Paniccia, "High-speed optical modulation based on carrier depletion in a silicon waveguide," *Opt. Express*, vol. 15, no. 2, pp. 660–668, Jan 2007. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-15-2-660>
- [9] D. J. Thomson, F. Y. Gardes, Y. Hu, G. Mashanovich, M. Fournier, P. Grosse, J.-M. Fedeli, and G. T. Reed, "High contrast 40gb/s optical modulation in silicon," *Opt. Express*, vol. 19, no. 12, pp. 11 507–11 516, Jun 2011. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-19-12-11507>
- [10] P. Dong, S. Chandrasekhar, X. Liu, L. L. Buhl, R. Aroca, Y. Baeyens, and Y.-K. Chen, "224-gb/s pdm-16-qam modulator and receiver based on silicon photonic integrated circuits." Optical Society of America, 2013.
- [11] P. Dong, L. Chen, C. Xie, L. L. Buhl, and Y.-K. Chen, "50-gb/s silicon quadrature phase-shift keying modulator," *Opt. Express*, vol. 20, no. 19, pp. 21 181–21 186, Sep 2012. [Online]. Available: <http://www.opticsexpress.org/abstract.cfm?URI=oe-20-19-21181>
- [12] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. Beausoleil, and J. Ahn, "Corona: System implications of emerging nanophotonic technology," in *Computer Architecture, 2008. ISCA '08. 35th International Symposium on*, June 2008, pp. 153–164.
- [13] G. Kurian, J. E. Miller, J. Psota, J. Eastep, J. Liu, J. Michel, L. C. Kimerling, and A. Agarwal, "Atac: A 1000-core cache-coherent processor with on-chip optical network," in *Proceedings of the 19th International Conference on Parallel Architectures and Compilation Techniques*, ser. PACT '10. New York, NY, USA: ACM, 2010, pp. 477–488. [Online]. Available: <http://doi.acm.org/10.1145/1854273.1854332>
- [14] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary, "Firefly: Illuminating future network-on-chip with nanophotonics," *SIGARCH Comput. Archit. News*, vol. 37, no. 3, pp. 429–440, Jun. 2009. [Online]. Available: <http://doi.acm.org/10.1145/1555815.1555808>
- [15] Y. Pan, J. Kim, and G. Memik, "Flexishare: Channel sharing for an energy-efficient nanophotonic crossbar," in *High Performance Computer Architecture (HPCA), 2010 IEEE 16th International Symposium on*, Jan 2010, pp. 1–12.
- [16] D. Vantrease, N. Binkert, R. Schreiber, and M. Lipasti, "Light speed arbitration and flow control for nanophotonic interconnects," in *Microarchitecture, 2009. MICRO-42. 42nd Annual IEEE/ACM International Symposium on*, Dec 2009, pp. 304–315.
- [17] A. Garca-Guirado, R. Fernandez-Pascual, J. M. Garca, and S. Bartolini, "Managing resources dynamically in hybrid photonic-electronic networks-on-chip," *Concurrency and Computation: Practice and Experience*, vol. 26, no. 15, pp. 2530–2550, 2014. [Online]. Available: <http://dx.doi.org/10.1002/cpe.3332>
- [18] Y. Li, Y. Zhang, L. Zhang, and A. W. Poon, "Silicon and hybrid silicon photonic devices for intra-datacenter applications: state of the art and perspectives

invited

, " *Photon. Res.*, vol. 3, no. 5, pp. B10–B27, Oct 2015. [Online]. Available: <http://www.osapublishing.org/prj/abstract.cfm?URI=prj-3-5-B10>

- [19] R. Ubal, B. Jang, P. Mistry, D. Schaa, and D. Kaeli, "Multi2sim: A simulation framework for cpu-gpu computing," in *PACT*. ACM, 2012, pp. 335–344.
- [20] G. Chen, H. Chen, M. Haurylau, N. Nelson, P. M. Fauchet, E. G. Friedman, and D. Albonese, "Predictions of cmos compatible on-chip optical interconnect," in *Proceedings of the 2005 International Workshop on System Level Interconnect Prediction*, ser. SLIP '05. New York, NY, USA: ACM, 2005, pp. 13–20. [Online]. Available: <http://doi.acm.org/10.1145/1053355.1053360>