

Contrasting Wavelength-Routed Optical NoC Topologies for Power-Efficient 3D-stacked Multicore Processors using Physical-Layer Analysis

Luca Ramini[†], Paolo Grani[§], Sandro Bartolini[§], Davide Bertozzi[†]

[†] ENDIF, University of Ferrara, 44122 Ferrara, Italy.

[§] Computer Engineering Department, University of Siena, 53100 Siena, Italy.

luca.ramini@unife.it[†], grani@dii.unisi.it[§], bartolini@dii.unisi.it[§], davide.bertozzi@unife.it[†]

Abstract—Optical networks-on-chip (ONoCs) are currently still in the concept stage, and would benefit from explorative studies capable of bridging the gap between abstract analysis frameworks and the constraints and challenges posed by the physical layer. This paper aims to go beyond the traditional comparison of wavelength-routed ONoC topologies based only on their abstract properties, and for the first time assesses their physical implementation efficiency in an homogeneous experimental setting of practical relevance. As a result, the paper can demonstrate the significant and different deviation of topology layouts from their logic schemes under the effect of placement constraints on the target system. This becomes then the preliminary step for the accurate characterization of technology-specific metrics such as the insertion loss critical path, and to derive the ultimate impact on power efficiency and feasibility of each design.

I. INTRODUCTION

One of the main drivers for considering optical interconnect technology for on-chip communication is the expected reduction in power. However, despite the arguments in favour of optical networks-on-chip (ONoCs) and the promising integration route, ONoCs are currently only at the stage of an appealing research concept. Understanding the implications of the specific properties of optical links across the upper layers of ONoC design is key to evolving ONoCs to a mature interconnect technology with practical relevance.

A fundamental decision in the early stage of ONoC design which may greatly benefit from this approach consists of topology selection. In fact, ONoC topologies are typically proposed in terms of their logic schemes, or are tied to specific floorplanning assumptions [3]. Therefore, the expected congruent multiples in communication performance or power savings may not materialize in practice.

On one hand, there might be a profound difference between the logic topology and its physical implementation [6], which raises the design predictability concern for ONoCs as well. Insertion loss, crosstalk and power analysis are important steps to tackle such a concern [2], and to assess the actual feasibility of connectivity patterns from a physical-layer standpoint.

On the other hand, a realistic assessment of topology implementation efficiency is not feasible if placement and routing constraints on the target system are not accounted for, which is a typically overlooked issue. This set of constraints strictly depends on the ultimate integration strategy of the optical interconnect with the electronic one. 3D integration today exhibits the capability to inexpensively integrate heterogeneous technologies while mitigating the compound yield risks. Therefore, it is reasonable to expect an optical layer stacked on top of an electronic one. However, the existence of interfaces between electronic and photonic signals implies strong constraints on the layout of the 3D architecture [10], that might break the regularity assumptions of ONoC connectivity patterns, or the floorplanning assumptions they are tied to.

The impact of place&route constraints might be especially severe for wavelength-routed ONoC topologies (WRONoC).

In fact, in WRONoCs the switching functionality is implemented using wavelength filters throughout the network. This implementation style removes control tasks from the critical path (conflict-free routing is guaranteed from the ground up by wavelength selection for each source-destination pair), enables predictive communication performance regardless of ongoing communications and does not require dynamically reconfigurable switching elements (via dual electronic NoCs). While appealing for low latency, WRONoCs share the full throughput that optics can provide among multiple communication flows, rather than devoting it entirely to a specific flow like in space-routed ONoCs. As a consequence, topologies have been mainly optimized to permanently provide full connectivity while minimizing the number of wavelengths and of physical resources. This has led to tightly optical technology-specific topologies ranging from rings [12] to customized multi-stage networks [7], [8], [9], which often make strong and unrealistic assumptions on master and slave placement or total wirelength to achieve compact and efficient implementation.

This paper targets the technology- and layout-aware characterization of relevant WRONoC topologies, thus aiming at more trustworthy comparative results than abstract comparison frameworks. For this purpose, the physical implementation efficiency of topologies under test is assessed in an homogeneous experimental setting with practical relevance, namely a 3D-stacked multicore processor with an optical layer targeting inter-cluster as well as processor-memory communication. Topologies will be compared in their ability to deliver the same communication bandwidth with the minimum power consumption. The novel contributions of this paper are:

- A.** A full custom place&route of multiple WRONoC topologies is performed, subject to the placement constraints of the target system. This way, the gap between logic topologies and their physical implementations is quantified in comparative terms.
- B.** The ultimate implications of physical properties on total power consumption are derived for each topology, thus quantifying the power gap between them (if any) and how to exploit it to increase wavelength parallelism.
- C.** A new WRONoC topology named *snake* is proposed, aiming at an implementation that better matches the placement constraint of the target system.
- D.** Switch-less optical rings will be compared with topologies relying on photonic switching elements (PSEs), thus assessing the actual need for these latter in the context of WRONoCs. The conclusion on this topic will be supported by preliminary scalability results on the same target system.
- E.** In order to increase the level of confidence of this comparative framework, we will not consider naive implementations of topologies, but optimization techniques of high practical relevance will be applied to them, such as spatial division multiplexing (for the ring), network partitioning for wavelength reuse (all topologies), and slight topology transformations for more flexible and/or efficient place&route (for the optical crossbar and GWOR topologies).

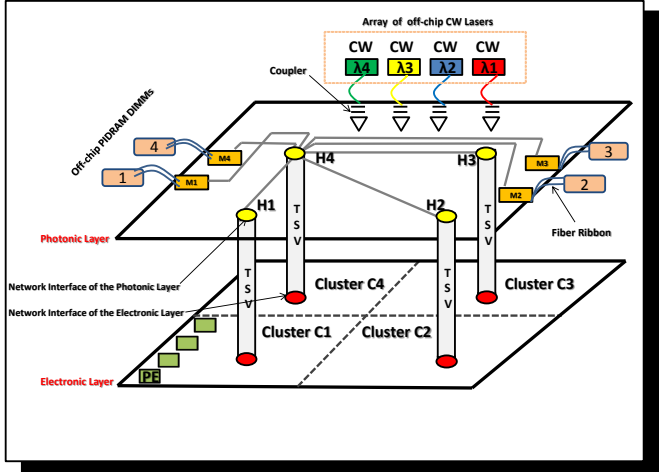


Fig. 1. Target 3D Architecture

II. 3D-TARGET ARCHITECTURE

The common experimental setting of practical interest to assess WRONoC topologies is a 3D architecture for multicore processors (see Fig.1), consisting of an electronic layer and of an optical one stacked on top of it. We assume that 64 identical processor cores are structured into 4 clusters of 16 cores, each cluster having its own gateway to the optical layer. We assume an area footprint of 1 mm^2 for each core, and a die size of $8 \text{ mm} \times 8 \text{ mm}$.

This latter is designed to accommodate three kinds of communications: (a) among clusters; (b) from a cluster to a memory controller of an off-chip photonically integrated DRAM DIMM [4]; c) from a memory controller to a cluster.

The optical layer is characterized by precise placement constraints imposed by the 3D-stacked architecture that topology layouts should satisfy. The first one consists of the position of the hubs. The aggregation factor (i.e., number of cores per cluster) and the total number of cores in the electronic plane dictate the position of the gateways and consequently of the optical network interfaces in the optical plane. As a consequence, we organize hubs along a square in the middle of the optical layer (see H1,H2,H3 and H4 in Figure 1).

In addition, we assume 4 memory controllers (M1,M2,M3 and M4) located pairwise at the opposite extremes of the chip, as proposed in conventional chip multiprocessor architectures, thus avoiding centralized communication bottlenecks for the on-chip network.

The above placement constraints radically question the practical feasibility of topology logic schemes and make the design of their associated real topology layout mandatory. In our system, we need to connect 8 initiators (4 hubs, 4 memory controllers) with 8 targets (the target interface of the same 4 hubs and 4 controllers). For this purpose, we revert to wavelength-routed optical NoCs, which allow contention-free communication and do not incur any path-setup/teardown overhead unlike space-route ONoCs [1], [2], [5]. WRONoCs deliver permanent full connectivity, i.e., all masters can potentially communicate with all slaves at the same time. The underlying principle is twofold: each master uses a different wavelength for each slave, and each slave receives packets from the different masters on different wavelengths. The interconnect fabric should avoid any interference between packets sent by different masters on the same wavelengths. Clearly, topologies with fewer physical resources will force the use of a higher number of wavelengths to enable conflict-free communication. The price that WRONoCs pay to deliver full connectivity consists of the serialization of a bit-parallel electronic flit onto

a destination-specific modulation wavelength, although some degree of broadband switching is feasible [6], [10].

This work does not blindly apply topologies under test to the master/slave connectivity problem of the target system, since the paper in [6] has demonstrated that even at such a small system scale a typical global topology for all communication actors is infeasible: too many waveguide crossings arise in an attempt to accommodate the connectivity pattern onto the 2D floorplan. As a consequence, [6] suggests the use of network partitioning, not only as a means of increasing design predictability, but also of enabling wavelength (and laser source) reuse across partitions. This work builds on the conclusions of [6] and takes the ONoC partitioning approach. In particular, we devote each network partition to a specific traffic class, namely inter-cluster communications, memory access requests from clusters and memory responses from memory controllers. A topology is mapped to each partition. However, this strategy enables to cut down on the number of wavelengths from 8 to just 4 due to their reuse.

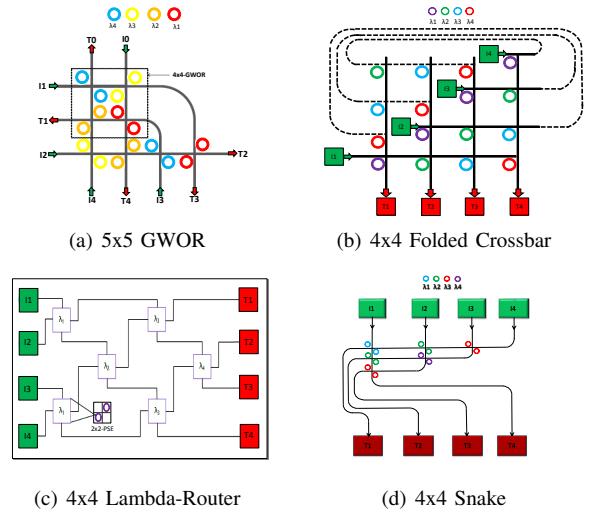


Fig. 2. Logic schemes of WRONoC topologies under test

III. LOGIC TOPOLOGIES

This section illustrates the logic scheme of WRONoC topologies under test, considering that each network partition will have to interconnect at most 4 masters with 4 slaves. We consider the most relevant schemes that have been proposed so far in the open literature, in addition to engineering an ad-hoc topology for the 3D-stacked system at hand.

[3] presents **4x4-GWOR**, a scalable and non-blocking wavelength-routed optical router. The basic cell is tied to a specific placement of actors, since it has 4 bidirectional ports located on the cardinal points. Two horizontal and two vertical waveguides are used, which intersect each other to form a basic check shape. MRRs (Micro-Ring-Resonators) are placed pairwise on waveguide intersections. GWOR does not support self-communication, hence its use for the memory request and response networks requires its extension to a 5x5 configuration. This is possible, since the wavelength assignment in [3] enables any size of the topology. As you can see in Fig.2(a), 5x5-GWOR is constructed starting from its lower basic cell (4x4-GWOR). With respect to the baseline scheme, we had to add 3 MRRs to work around the lack of self-communication and enable each master to be connected with 4 slaves. At the same time, one input is unused, therefore redundant MRRs were removed.

An alternative topology is illustrated in [7] and is named **4x4-lambda Router**. In order to interconnect 4 masters with

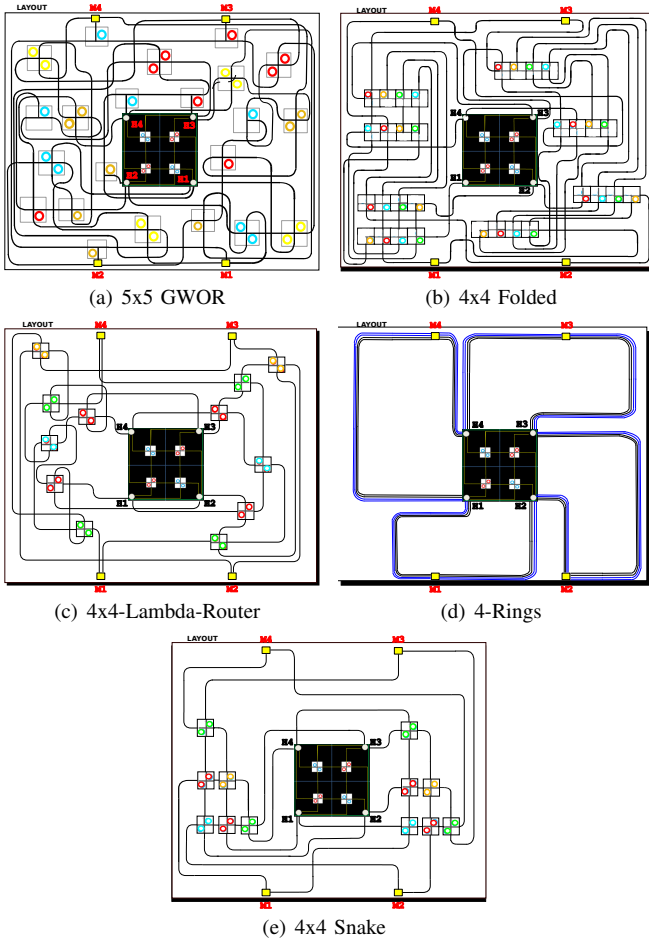


Fig. 3. Layout of the Optical layer with network partitioning after manual place&route. Requests networks are on the left while response ones on the right of the layout.

4 slaves, the network makes use of 4 stages of 2 and 1 add-drop optical filters (Fig.2(c)). The topology resembles that of electronic multistage interconnection networks, although the connectivity pattern is strictly customized for the optical technology, and for the needs of wavelength routing in particular. With respect to the original scheme, we replaced the native 2x2 add-drop filters with 2x2 photonic switching elements, the only difference being an easier physical design thanks to the orthogonally intersected waveguides.

As illustrated in Fig.2(b), an optimized optical crossbar, here referred to as **4x4 Folded Crossbar**, was customized for connecting 4 initiators with 4 targets. With respect to other solutions, the logic scheme of this topology makes use of long optical links to interconnect all communication actors and only embeds 1x2-PSEs, hence potentially resulting in the largest number of MRRs. With respect to the standard scheme of the crossbar, we counterintuitively misaligned the injection points of masters, thus causing the need for wrap-around links. However, this is only an illusory effect of the logic scheme, since this optimization gives more flexibility to the physical design of the topology and the total wire length in the layout is actually shorter than for the standard crossbar.

Le Beux et al. developed an Optical Ring topology in [12], called **ORNoC**, together with its optimized wavelength assignment policy. A single wavelength is reused for multiple parallel communications across the same waveguide by avoiding their overlapping. This way, scalability is facilitated while containing the number of physical waveguides. The key property of ORNoC is that in principle it has neither

waveguide crossings nor photonic switching elements, which makes it an appealing solution with respect to those reported so far. However, there are key effects that come into play when actual implementation is pursued. First, the amount of physical resources is so small that conflict-free wavelength routing becomes infeasible on a single waveguide even for the small scale system targeted by this paper, unless a large number of wavelengths and laser sources is used. Therefore, this paper takes the use of spatial division multiplexing for granted for optical rings, i.e., communications are spread across multiple physical waveguides. Second, reachability of all waveguides from masters and slaves cannot avoid undesired crossings even in a 3D-stacked scenario. At least, light is modulated on the optical layer, and should then reach even the waveguides that are further away from the modulation point. The receiver part can be instead optimized, since photodetector outputs could go directly into the electronic plane through TSVs without crossing any waveguide. Third, MRRs are anyway needed to inject wavelengths into and extract them out of the waveguides. Fourth, for large chips, the propagation loss of the long ring waveguides becomes significant and is certainly the major contributor to the insertion loss of this topology. All together, it is not clear whether the above inconveniences can offset the theoretical benefits of rings with respect to switch-rich and crossing-prone topologies. This paper sheds light on this issue in the context of WRONoCs, where topologies have to deliver the same bandwidth and comparable latency. For the sake of comparison, we will constrain all topologies to use the same number of wavelengths and laser sources, and to instantiate physical resources accordingly.

Finally, in this paper we propose a novel scalable and contention-free logic scheme, named the **Snake topology**. The pattern (Fig.2(d)) is also flexible, since a different number of initiators and targets can be easily accommodated. In the 4x4-Snake, six wavelength filters (2x2-PSEs) are tuned to different wavelengths and their number scales up from the rightmost side to the leftmost one. 4 main optical links have a winding shape and connect the slaves while enabling some placement flexibility. This topology was conceived to map efficiently to the placement constraints of the target system, and should be viewed as a custom-tailored solution for the system at hand.

IV. PHYSICAL TOPOLOGIES

This section deals with the problem in assigning topologies to network partitions and to lay them out. For the inter-cluster ONoC, the choice is trivial: 4x4-GWOR delivers the needed connectivity in a scenario where its physical placement assumptions are perfectly satisfied. At the same time, it features the lowest number of MRRs. Therefore, we restrict the problem of identifying the topologies that are better suited for processor-memory communication, and lay them out twice: for the memory request network (from hubs to memory controllers) and the memory response one (from controllers to hubs). The fundamental difference lies in the flipped position of masters and slaves, which makes them asymmetric.

Due to the lack of automatic place&route tools for optical NoCs, we manually placed and routed the topologies, hence coming up with full custom design solutions. We only did not consider the routing of the light distribution network. The methodology and the design rules adopted for the physical implementation of each logic topology were inspired by those used for multi-stage electronic networks like fat-trees [11]. First, each switch is placed close to its attached node; second, switches without any node connection are homogeneously spread across the floorplan in order to balance length of waveguides, and above all to avoid waveguide crossings. Since these latter play a dominant role in determining the minimum optical power that laser sources should provide to satisfy specific detector sensitivities, we consider two relevant and

TABLE I
LAYOUT-AWARE PROPERTIES OF TOPOLOGIES UNDER TEST

	Total number of Wavelengths	Max number Crossing	Max wire length cm	Total number of MRRs	Type of MRR
4-RINGS	4	3	3.2	40 (8 IC)	4
4x4 SNAKE	4	6	2.4	32 (8 IC)	4
4x4 λ -Router	4	15	1.8	32 (8 IC)	4
4x4 Folded Crossbar	4	21	2	40 (8 IC)	4
5x5 GWOR	4	31	2.4	40 (8 IC)	4

increasingly aggressive optimizations: elliptical tapers [13] and Multi-Mode-Interference (MMI) tapers [14].

In spite of these efforts, the difference between logic and physical topologies is still apparent. In some cases, waveguides become circuitous and additional waveguide crossings cannot be avoided unlike the small system scale, mitigated only by the use of network partitioning. 5x5-GWOR (Fig.3(a)) suffers from the different placement position of network interfaces with respect to the logic scheme, to such an extent that the critical path increases from 4 crossings to 31. Despite a higher worst case number of crossings in the logic scheme (6), the layout of the 4x4 Folded Crossbar Fig.3(b) resulted only in 21 crossings, with the same number of MRRs.


The layouts of the 4x4-lambda Router (Fig.3(c)), ORNoC (Fig.3(d)), and 4x4-Snake (Fig.3(e)). are clearly less intricate than the previous ones, hence potentially resulting in lower insertion loss critical paths. More precisely, Lambda-Router counts 15 crossings while Snake only 6. By using the wavelength assignment in [12] and a convenient ordering of nodes along waveguides, ORNoC turns out to exhibit 3 crossings on the critical path, all localized close to network interfaces for the sake of waveguide reachability. This represents a significant optimization. Key properties of topologies under test, measured after their physical design, are summarized in table.I. They are referred to the network as whole, inclusive of the three partitions. While all topologies natively used 4 wavelengths, a spatial division multiplexing over 4 waveguides had to be used for ORNoC to achieve the same goal.

Surprisingly, Snake and Lambda-Router solutions make use of 32 MRRs (24 in the request and response networks vs. 8 in the inter-cluster one) against 40 of the Ring one. The key reason lies in the fact that each optical network interface in the ring needs 4 MRRs to inject modulated wavelengths into their waveguides, in addition to 8 rings needed in the inter-cluster network. All other topologies instead do not have any injection filters, since they get a branch of the light distribution network which directly enters the network. In the ring, the injection waveguide needs to be bridged to the ring waveguides. Extraction filters at receivers are common for all topologies, hence were not considered in the count.

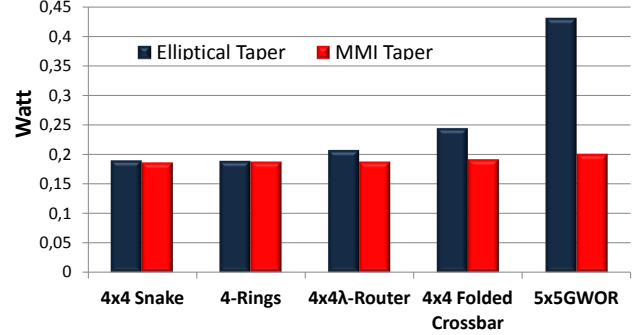
V. EXPERIMENTAL RESULTS

As a photonic message propagates through the network, it is attenuated by multiple physical contributions such as waveguide scattering, ring resonator loss, and waveguide crossing reflections, that build up the breakdown of the total network-level insertion loss.

For this purpose, we first quantify the critical path insertion loss IL_{max} of all multi-partition topologies investigated so far. Once IL_{max} is obtained and the detector sensitivity is known (e.g. $S = -17dBm$ [15]), it is possible to determine the lower limit of optical laser power (P) to reliably detect the corresponding photonic message at the destination node. We quantify the worst case IL_{max} on each wavelength across all partitions and we consequently derive the global topology IL_{max} . We then make the practical assumption that such a



(a) Max Insertion Loss across topologies



(b) Total Power across topologies

Fig. 4. Power-efficiency of topologies under test

TABLE II
PARAMETERS USED IN THIS WORK

Parameters	Value	Devices	Features
Propagation-Loss [2]	1.5dB/cm	Laser	CW(Continuous Wave) PLE=20% (Laser efficiency)
Bending-Loss[2]	0.005dB		PCW=90% (Coupling Laser-Link)
Crossing-Loss		Modulator	Si Disk $\beta = 20\%$ (Launch efficiency)
Optimized by Elliptical Taper[22]	0.52dB		Dyn. Dissipation=3fJ/bit Static Power=30W Vdd=1V Modulator Power depends on IL_{max} [16]
Optimized by MMI Taper[22]	0.18dB		
Drop-Loss		Detector	CMOS(45nm) hybrid silicon receiver S=-17dBm, (BER= 10^{-12} @10Gbit/s) Power=3.95mW [15]
Optimized by Elliptical Taper[22]	0.013dB		
Optimized by MMI Taper[22]	0.0087dB	Photonic-Switching-Elements (PSEs)	Thermal-Tuning: 20 μ W/ring [2]

worst case IL_{max} dictates the power requirement for all laser sources.

Our study assumes loss parameters reported in table.II. We rely on a Simulink simulation framework to quantify physical metrics of optical networks. We first simulate every single path of a specific topology taking into account the above loss parameters; then, we calculate the corresponding insertion loss as the sum of all components (PSEs, straight, bend and crossing waveguides and drop-into-ring losses) which affect the path under test. The topology models assume die sizes of 8 mm x 8 mm.

A. Power efficiency of topologies

Figure 4(a) shows the worst-case insertion loss across all topologies considered in this comparison, with both kinds of tapers at waveguide crossings. GWOR turns out to be the worst solution, since it suffers from 31 crossings and 24 mm of wiring length on the critical path while ORNoC (the best

solution) has just 3 crossings but 32 mm of waveguides. The Snake topology, with its 6 crossings and the same max length of GWOR, becomes competitive, since propagation losses are still not very relevant at this chip size. With elliptical taper, the overhead with respect to ORNoC is just 5%. 4x4-Lambda Router has reasonable results in the comparison since it has 18mm of wiring length and 15 crossings, while the 4x4-Folded Crossbar is better than GWOR for two reasons: lower number of crossings (21), and 4 mm shorter link length.

The effect of MMI is highly beneficial for the Snake, since it minimizes the impact of its crossings over IL_{max} , while benefits are not so relevant for the waveguide-dominated ORNoC. This latter ends up in a 13.2% higher insertion loss than Snake. This result is very interesting, since it points out that there is actually a role also for non-ring topologies in WRONoCs, in spite of their apparent higher complexity. On the other hand, the feasibility of MMI should not be taken for granted, since it depends on the maturity of the manufacturing process and on the device size. In turn, Snake results in a 13.8%, 32.6% and 49.5% lower insertion loss than Lambda-Router, Folded Crossbar and GWOR respectively.

By using such critical path insertion losses, it was possible to derive the needed laser power to meet a bit-error-rate (BER) [17] of 10^{-12} at the optical receivers with a fixed sensitivity of -17dBm [15]. It was then possible to account for the power contribution of modulators [16], detectors [15], and thermal tuning [2], thus estimating total power for each topology. Relevant parameters are in table.II.

Figure 4(b) shows the total power across all topologies when the energy consumption of the detector is 395fj/bit (or 3.95mW), as demonstrated in [15]. Power refers to the scenario where the maximum aggregate bandwidth of the network is used (around 440Gbit/sec with modulation rates of 10Gbit/sec). As you can see, the total power of GWOR is higher than that of other topologies regardless of specific taper. With elliptical tapers, GWOR is clearly infeasible under the given place&route constraints, and so is the folded crossbar. The capability of the Snake topology to track power efficiency of the optical ring (the best solution) is remarkable at this system scale.

The effect of MMI tapers is to reduce the critical path differentiation across topologies, hence significantly bridging the gap between the best and the worst one. Laser and modulator power are closely related to the IL_{max} of the topologies, however the total network power is dominated by receiver power with current technology assumptions (it counts on average 75% with Elliptical taper while 90% with MMI taper), therefore the remaining gap between topologies in Figure 4(a) maps to the total power gap of Figure 4(b) after going through an attenuation factor: just 15mW of difference between Snake (the best) and GWOR (the worst). Of course, different laser source (e.g., efficiency) or receiver (e.g., energy) parameters may further widen again the gap.

As a next step, we want to characterize the impact of system scale and technology evolution on this trend. For this purpose, we sketch a future generation of the target system. We now assume 128 cores in the tile-based electronic plane, getting access to the optical layer through 8 gateways (and 8 corresponding hubs in the optical plane). The number of memory controllers is kept the same, which might be possible due to the benefits of photonic integration deeper into the DRAM DIMM [4]. Consequently, the die sizes grow to 16 mm x 16 mm. We limit the comparison between ORNoC and the best topology found so far, i.e., the Snake, and omit the inter-cluster network. Therefore, we manually placed and routed two 4-waveguide ORNoCs and two separate Snake topologies (an asymmetric 8x4 for memory requests and a 4x8 to enable memory responses). We assume MMI tapers to be mainstream in these topologies and that detector energy can be improved up to 50fj/bit [2] while conservatively keeping




Fig. 5. Max Insertion-Loss under Scaled Assumptions Contrasting Snake vs. Rings




Fig. 6. Total Power under Scaled Assumptions Contrasting Snake vs. Rings

the same sensitivity, a projection which is supported by the physical considerations in [18] about silicon photonics in 3D-stacked systems and receiver circuitry.

Figure.5 shows the insertion loss critical path breakdown of each topology. The 8 rings are in fact heavily penalized by the high wiring length over the new die size (64 mm vs. 48 mm of Snake), which leads to a larger amount of propagation loss regardless of the higher number of crossing losses in Snake (1.75x higher than 8-Rings).

The total power consumption across the two topologies is shown in Figure.6. Thanks to the lower insertion-loss on the critical path and the higher maturity of receiver technology, Snake results more efficient than ORNoC by about 15%. This certainly confirms that optical rings are not the most power efficient and least complex solution under all WRONoC scenarios, although conclusions are tightly instance- and technology-specific.

B. System-Level Implications

In section V-A we pointed out a significant power gap between GWOR and ORNoC (or Snake) in the target system in the presence of crossings optimized with elliptical tapers. In this section we show that the most power efficient topologies might use this power budget (around 250mW) to increase their wavelength parallelism. This would decrease the serialization ratio at the electro-optical network interface and improve system performance. This is typically referred to as broadband switching. We computed that a 250mW gap would enable ORNoC/Snake a wavelength parallelism of 2 on every master-slave optical channel, including the cost for the additional modulators and receivers. This would mean around

TABLE III
PARAMETERS OF THE SIMULATED ARCHITECTURE

Cores	4 clusters, 1 GHz cores
L1 caches	16 kB + 16 kB Instr./Data, 4-way, 1 cycle hit time
L2 cache	4 MB, 8-way, shared and distributed 16x256 kB banks, 2/5 cycles tag/tag+data (bank)
Coherency	MOESI, distributed directory and one per cluster memory controller
NoC	Electronic mesh intra cluster, 32 bit, 1 GHz WRONoC inter-cluster and processor-memory, 1/2/4 bit
Main memory	1 GByte, DDR2 DRAM, 80 cycles




Fig. 7. System-level performance speedup (normalized).

80Gbit/sec of memory traffic from each hub. Alternatively, the wavelength budget might be allocated heterogeneously across the channels, devoting more bandwidth to the most congested ones. To quantify this benefit, we performed a system-level simulation where we implemented these features.

Full system evaluation was obtained using the gem5 simulator [19], in which we model the clustered 16-core architecture described in Table III and employing our WRONoC partitions for inter-cluster communication as well as for communication towards and from main memory through four memory controllers. Simple local NoCs are used for intra-cluster communication. Cache parameters were derived from Cacti 6.0 [21]. Performance were evaluated for the Parsec 2.1 multithreaded benchmark suite [20], which encompasses heterogeneous real-world applications for which we have used the medium input set. Linux 2.6.27 operating system (OS) was booted on the simulated architecture and we enforced core-affinity to reduce OS scheduling effects in successive runs.

Figure 7 shows the performance improvements that can be achieved at system level when different degrees of broadband switching are used and under the load of real-world complex benchmarks. We assume that the wavelength budget is homogeneously spread across all optical channels. In particular, 2-bit parallelism (the case of interest) allows for more than 52% average improvement and up to 61% for *bodytrack* application, while 4-bit parallelism reaches 68% average improvement with a peak of 80% for *canneal*.

Using more than 4-bit optical parallelism is useless as performance saturates by construction. In fact, the proposed contention-free network topology allows concurrent optical communications between each core pair without contention and with the indicated parallelism. As each electronic link towards the optical path feeds the electro/optical hub at 32 Gbps (32bit/flit @ 1GHz), a 4-bit optical interface working at 40Gbps is able to drain the communication at full speed without inducing any queuing. Therefore, a wider optical interface would be idle for most of its time and could not be able to improve communication performance in any way. Removing such interface bottleneck is outside the scope of this paper.

These results highlight that part or all of the power saved by ORNoC or Snake over GWOR can be fruitfully used to improve overall system performance and still maintaining a power advantage over the baseline.

VI. CONCLUSION

In this paper, we performed a comparative analysis of WRONoC topologies by considering both the properties of optical links as well as placement constraints on a target system of practical interest. With elliptical tapers, already at small system scales, some topologies are impractical and a large power gap does exist, which could be exploited

for performance-efficient broadband switching. At the same time, optical rings and customized switching networks provide roughly the same power, although rings are simpler. However, in those application scenarios where connectivity requirements and die size increase, spatial division multiplexing combined with the relevant role of propagation losses seriously penalizes optical rings. Even for small scale scenarios, should technology evolutions improve optical receiver energy, switching networks could again have a role. In practice, an optical ring is ideally the best WRONoC topology, but its practical non-idealities (e.g., waveguide reachability, injection system, worse waveguide length scalability) make an actual comparative test with other topologies mandatory in the target system.

A key takeaway is however that abstract or even pencil-and-paper floorplanning considerations might lead to misleading comparative results. This makes the case for the development of automatic place&route tools, which we will pursue in our future work.

VII. ACKNOWLEDGMENTS

This work was supported by PHOTONICA project (RBF08LE6V) under the FIRB 2008 program, funded by the italian government.

REFERENCES

- [1] A. Shacham, K. Bergman, L. P. Carloni, "On the Design of a Photonic Network-on-chip", NOCS'07: International Symposium on Networks-on-Chip, May 2007.
- [2] J. Chan et al., "Architectural Exploration of Chip-Scale Photonic Interconnection Network Designs Using Physical-Layer Analysis", Journal of Lightwave Technology, vol.28, n.9, pp.1305-1315, May 2009.
- [3] X. Tan et al., "On a Scalable, Non-Blocking Optical Router for Photonic Networks-on-Chip Designs", Photonics and Optoelectronics (SOPO), May 2011.
- [4] S. Beamer et al., "Re-Architecting DRAM Memory Systems with Monolithically Integrated Silicon Photonics", ISCA'10: International Symposium on Computer Architecture, June 2010.
- [5] A. Shacham, K. Bergman, and L. P. Carloni "Photonic Networks-on-Chip for Future Generations of Chip Multiprocessors", IEEE Trans. on Computers, vol.57, n.9, pp. 1246-1260, September 2008.
- [6] 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", NOCS'12: International Symposium on Networks-on-Chip, May 2012.
- [7] I. O'Connor et al., "Towards Reconfigurable Optical Networks on Chip", ReCoSoC 2005, pp.121-128.
- [8] A. Scandurra and I.O'Connor, "Scalable CMOS-compatible photonic routing topologies for versatile networks on chip", Network on Chip Architecture, 2008.
- [9] S. Le Beux et al., "Multi-Optical Network-on-Chip for Large Scale MP-SoC", IEEE embedded systems letters, vol.2, n.3, pp. 77-80, September 2010.
- [10] S. Le Beux, J.Trajkovic, I.O'Connor and G.Nicolescu, "Layout Guidelines for 3D Architectures including Optical Ring Network-on-Chip (ORNoC)", VLSI-Soc'11: International Conference on VLSI and System-on-Chip, October 2011.
- [11] D. Ludovici et al., "Assessing Fat-Tree Topologies for Regular Network-on-Chip Design under Nanoscale Technology Constraints", DATE'09: Conference on Design, Automation and Test in Europe, April 2009.
- [12] S. Le Beux et al., "Optical Ring Network-on-Chip (ORNoC): Architecture and Design Methodology", DATE'11: Conference on Design, Automation and Test in Europe, March 2011.
- [13] N. Sherwood-Droz et al., "Optical 4x4 hitless silicon router for optical Networks-on-Chip (NoC)", Opt. Expr., vol. 16, n. 20, pp. 15915-15922, 2008.
- [14] H. Chen and A.W Poon, "Low-Loss Multimode-Interference-Based Crossings for Silicon Wire Waveguides", Photonics Technology Letters, IEEE, vol. 18, n. 21, pp. 2260-2262, 2006.
- [15] Xuezhe Zheng et al., "Ultra-efficient 10Gbit/s hybrid integrated silicon photonic transmitter and receiver", Opt Express, 14;19(6):5172-86, March 2011.
- [16] David A.B. Miller, "Energy consumption in optical modulators for interconnects", Opt Express, Vol. 20, pp. A293-A308, March 2012.
- [17] G.P. Agrawal, "Fiber-Optic Communication Systems", Wiley-Interscience, third edition, chapter fourth, pp.133-178, 2002.
- [18] M.Georgas et al., "A Monolithically-Integrated Optical Receiver in Standard 45-nm SOI", Solid State Circuits, 2002.
- [19] L.Nathan et al., "The M5-Simulator: Modeling Networked Systems", IEEE MICRO, 2006.
- [20] C.Bienia et al., "The PARSEC Benchmark Suite: Characterization and Architectural Implication", PACT, 2008.
- [21] Naveen Moralimanoohar and Rajeev Balasubramonian, "CACTI 6.0: A toll to model large caches", IEEE MICRO, 2006.
- [22] G.R.Hadley, "Effective index model for vertical-cavity surface-emitting lasers", Opt.Lett., vol.20, pp.1483-1485, 1995.