

UNIVERSITY OF FERRARA

---

Engineering Department of the University of Ferrara

Doctorate Degree in Science of Engineering

Coordinator: Prof. Stefano Trillo

Cycle: XXVI

**Towards Compelling Cases  
for the Viability of  
Silicon-Nanophotonic  
Technology  
in Future Many-core Systems**

ING-INF/01

Candidate:

Luca Ramini

Advisor:

Prof. Davide Bertozzi

---

Academic Year 2011/2014



I would like to dedicate this contribution  
to my future wife Lisa and our future child



# Contents

<b>Thesis Abstract</b>	<b>15</b>
<b>Thesis Introduction</b>	<b>17</b>
<b>1 On-Chip Optical Communications</b>	<b>21</b>
1 Penetration of optical links into communications . . . . .	21
2 On-Chip Optical Communication: Why? . . . . .	24
3 Silicon Photonics as a Technology Enabler . . . . .	24
3.1 Optical Links . . . . .	25
3.2 Modulators . . . . .	26
3.3 Photonic-Switching Elements and Optical Routers for Optical Networks-on-Chip . . . . .	27
3.4 Photodetectors . . . . .	28
3.5 Laser Sources . . . . .	29
3.6 3D-Stacked integrated systems (3D-ICs) . . . . .	29
4 Conclusion . . . . .	30
<b>2 Optical Networks-on-Chip</b>	<b>31</b>
1 Introduction . . . . .	31
2 Optical Networks-on-Chip . . . . .	31
3 Space-Routed Optical Networks-on-Chip . . . . .	33
3.1 4x4 Torus . . . . .	33
3.2 4x4 Torus NX . . . . .	34
3.3 4x4 Square Root . . . . .	34
3.4 The Optical Fat-Tree: FONoC . . . . .	35
4 Wavelength-Routed Optical Networks-on-Chip . . . . .	36

5	Conclusion . . . . .	38
<b>3</b>	<b>Towards Trustworthy Crossbenchmarking Framework between ONoC and ENoC: The Golden Rules</b>	<b>39</b>
1	Pathfinding Requirements . . . . .	39
2	An Overview of the Golden Rules . . . . .	41
3	Conclusion . . . . .	44
<b>4</b>	<b>The Design Predictability Gap in Optical Networks-on-Chip Design</b>	<b>45</b>
1	Introduction . . . . .	46
2	3D-Target Architecture . . . . .	47
3	Electro/Optical & Opto/Electrical Network Interfaces . . . . .	50
4	Design Predictability Gap: Logic Scheme vs. Physical Layout . . . . .	52
5	Topology Exploration: Global Connectivity . . . . .	52
5.1	Relative Topology Comparison . . . . .	53
5.2	Physical Layer Analysis . . . . .	57
	(a) Insertion loss Analysis . . . . .	57
	(b) Power Analysis . . . . .	60
5.3	Comparison with an Optical Ring Topology . . . . .	61
6	Network Partitioning . . . . .	65
6.1	Logic Topologies . . . . .	65
7	Snake vs. <b>Lambda Router</b> . . . . .	67
8	Physical Topologies . . . . .	68
8.1	Power efficiency of topologies . . . . .	72
9	Global Connectivity vs. Partitioning . . . . .	74
10	Scalability Implications . . . . .	76
10.1	System-Level Implications . . . . .	77
11	Conclusion . . . . .	80
<b>5</b>	<b>Network-Interface Architecture for Wavelength-Routed Op- tical NoC Topologies</b>	<b>83</b>
1	Abstract . . . . .	83

2	Network Interface Architecture:	
	A More Detailed View . . . . .	84
2.1	Wavelength Routed NoC . . . . .	84
2.2	Message Dependent Deadlock Avoidance . . . . .	85
2.3	Buffering Sources . . . . .	86
2.4	Serialization and Deserialization Procedure . . . . .	87
2.5	Resynchronisation:	
	Source Synchronous Communication . . . . .	87
2.6	Backpressure Mechanism:	
	The case of the Credit-based Flow Control . . . . .	88
2.7	E/O and O/E Conversions . . . . .	88
3	Evaluation . . . . .	89
3.1	Methodology . . . . .	90
3.2	Latency Breakdown . . . . .	91
3.3	Transaction Latency . . . . .	91
3.4	Static Power & Energy-per-Bit . . . . .	93
4	Conclusion . . . . .	94
<b>6</b>	<b>Crossbenchmarking Framework between the Most Efficient ONoC and its Aggressive Electrical Baseline</b>	<b>97</b>
1	Abstract . . . . .	97
2	Introduction . . . . .	98
3	Target System . . . . .	99
4	Baseline Electronic NoC . . . . .	100
5	Wavelength-Routed Optical Ring Design . . . . .	101
5.1	Design Methodology . . . . .	102
5.2	The Waveguide Crossings Concern in Optical Ring Design . . . . .	105
5.3	Laser Power Assessment . . . . .	105
6	Power Modeling . . . . .	106
7	Experimental results . . . . .	108
7.1	Methodology . . . . .	110
7.2	Result discussion . . . . .	113
7.3	Systeml-Level Energy and Conclusion . . . . .	116

<b>7</b>	<b>CAD Support for Design and Validation of Optical Networks-on-Chip</b>	<b>119</b>
1	Why an Automatic Place&Route Tool for ONoC Design is needed? . . . . .	119
2	Introduction . . . . .	120
3	PROTON's properties . . . . .	122
3.1	Topology Specification Format . . . . .	122
3.2	Placement & Routing Algorithm . . . . .	123
4	Maximum Insertion Loss . . . . .	125
5	PROTON at work . . . . .	126
5.1	Manual Design vs. PROTON . . . . .	126
5.2	Best Topology Selection . . . . .	128
5.3	Scalability . . . . .	130
6	Conclusion . . . . .	131
<b>8</b>	<b>Network-Level Simulation Frameworks for Optical Networks-on-Chip</b>	<b>133</b>
1	Abstract . . . . .	133
2	Background & Motivations . . . . .	134
3	Technology-Aware SystemC Simulation for Optical Networks-on-Chip . . . . .	137
4	S-parameters modelling of a 1x2 PSE . . . . .	138
5	SystemC Modelling of a 4x4 Optical Switch . . . . .	140
6	SystemC Modelling of a 4x4 Square Root Topology . . . . .	144
7	Simulink Simulation Framework for Optical NoCs . . . . .	148
8	Conclusion . . . . .	149
	<b>Thesis Conclusions</b>	<b>151</b>
	<b>References</b>	<b>153</b>
	<b>Publications</b>	<b>163</b>
	<b>Acknowledgements</b>	<b>165</b>

# List of Figures

1.1	Penetration of optical links into communications . . . . .	22
1.2	Schematic of an Electronic Network-on-Chip (ENoC). . . . .	23
1.3	Switching Elements and Routers used to build Optical Networks-on-Chip. . . . .	27
2.1	Milestones regarding the Space Routed Optical Networks-on-Chip. . . . .	32
2.2	The Optical Fat-Tree: FONoC. . . . .	35
2.3	Milestones regarding the Wavelength Routed Optical Networks-on-Chip. . . . .	37
3.1	The Pathfinding Requirement . . . . .	40
4.1	3D-Target Architecture . . . . .	48
4.2	Electronic Network Interface: Transmission Side . . . . .	49
4.3	Array of modulators in the optical layer . . . . .	50
4.4	Array of filters and detectors in the optical layer . . . . .	51
4.5	Electronic Network Interface: Reception Side . . . . .	51
4.6	The predictability gap : (a) 8x8 $\lambda Router$ logic scheme, (b) 8x8 $\lambda Router$ Real Layout manually generated . . . . .	53
4.7	(a) 8x8 GWOR logic scheme, (b) 8x8 GWOR physical layout, (c) 8x8 Folded Crossbar logic scheme, (d) 8x8 Folded Crossbar physical layout . . . . .	54
4.8	Calculation of Insertion loss for a small Network Segment . . . . .	57
4.9	ILmax Contrasting: Physical Layout vs. Logic Scheme . . . . .	60
4.10	Total power Contrasting: Physical Layout vs. Logic Scheme . . . . .	61
4.11	Optical Ring Physical Layout . . . . .	62

4.12	Hub architecture of an Optical Ring with physical awareness . . . . .	63
4.13	ILmax Contrasting: 7-way Ring vs. 8x8 Folded Crossbar . . . . .	64
4.14	Total power Contrasting: 7-way Ring vs. 8x8 Folded Crossbar . . . . .	64
4.15	Logic schemes of WRONoC topologies under test . . . . .	66
4.16	Asymmetric 8x4 Snake . . . . .	69
4.17	Snake Properties . . . . .	70
4.18	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. . . . .	71
4.19	Contrasting of the maximum insertion loss across topologies . . . . .	72
4.20	Total power comparison across topologies . . . . .	72
4.21	Total power comparison: Partitioned vs, Global Ring . . . . .	75
4.22	Total power comparison: Partitioned vs, Global Ring . . . . .	75
4.23	ILmax under Scaled Assumptions: Snake vs. Rings . . . . .	77
4.24	Total Power under Scaled Assumptions: Snake vs. Rings . . . . .	78
4.25	System-level performance speedup (normalized). . . . .	79
5.1	Optical Network Interface Architecture. . . . .	84
5.2	Principle of the Wavelength Selective Routing. . . . .	85
5.3	Dependence between a Request and Response at the NI. . . . .	86
5.4	Latency breakdown of the optical NI with 3-bit parallelism and the optical Ring. . . . .	90
5.5	Latency of the most common communication patterns. For the ENOC, we include minimum, maximum, and average paths. . . . .	91
5.6	Static power and Energy-per-Bit of the NIs and the electronic a optical NoCs. . . . .	94
6.1	Tile 16 (from Tiler Corporation). . . . .	99
6.2	Principle of the designed Optical Ring Architecture. . . . .	101
6.3	Floorplan of a 16x16 Wavelength-Routed Optical Ring Architecture. . . . .	103
6.4	2x2 Optical Ring Topology with physical awareness. . . . .	104
6.5	Laser power results across wavelengths: aggressive vs. realistic. . . . .	107
6.6	Performance comparison of the ONoC with the electronic baseline. . . . .	112

6.7	Energy comparison of the 3 bit (2nd bars) and 4 bit (3rd bars) ONoC wth respect to ENoC baseline for the common aggressive.	112
6.8	Energy comparison of the 3 bit (2nd bars) and 4 bit (3rd bars) ONoC wth respect to ENoC baseline for the common realistic.	113
6.9	Energy comparison of the 3 bit (2nd bars) and 4 bit (3rd bars) ONoC wth respect to ENoC baseline for the accurate realistic.	114
6.10	Energy comparison of the 3 bit (2nd bars) and 4 bit (3rd bars) ONoC wth respect to ENoC baseline for the accurate aggressive.	115
6.11	Normalized System-Level Energy Comparison. . . . .	116
7.1	8x8 $\lambda Router$ logic scheme. . . . .	123
7.2	An example of optical paths. . . . .	124
7.3	Propagation and Crossing Loss are tightly interrelated. . . . .	125
7.4	ILmax (dB) Contrasting: Manual vs. PROTON. . . . .	126
7.5	Laser Power (Watts) Contrasting: Manual vs. PROTON. . . . .	127
7.6	Table I: Results for variation of propagation and crossing weights.	129
7.7	16x16 $\lambda Router$ under scalability assumptions. . . . .	130
7.8	Results under scalability assumptions. . . . .	131
8.1	(A-left) Sketch of a microring resonator with orthogonal access waveguides. (A-right) Sketch of an orthogonal planar crossing.(B) The 1x2 PSE can be considered as the cascading of the microring the crossing. If the wavelength of the optical carrier signal is resonant with the microring, the data stream on the Port 1' is routed toward Port 2'; if the optical signal is not resonant, the data stream is routed to Port 3'. Dually, a signal coming from Port 4' is routed to Port 3' if resonant, while continuing to Port 2' if out of resonance. The device is bidirectional, i.e. each of the ports can act as input or output, but not simultaneously. (C) Representation of the PSE as it appears when discretized inside the 2D-FDTD code; in this case the transmission at the intersection between the waveguides is optimized using a MMI-based crossing. . . . .	139

---

8.2	Comparisons between the spectral response of the Through (Drop) path [Port 1' to Port 3'] ([Port 1' to Port 2']) of the passive 1x2 PSE, calculated via 2D-FDTD and by means of the s-matrix mode . . . . .	141
8.3	Topological scheme of the 4x4 Optical Switch showing the interconnections between the eighth PSEs. (B) FDTD representation of the real device. Down, Simulated (solid line) and modeled (dashed line) transmission curves for the paths linking the I-WEST port with the O-NORTH port and I-WEST port with the O-SOUTH port of the 4x4 Optical Switch. . . .	142
8.4	4x4 Square Root Topology. . . . .	145
8.5	Insertion-Loss comparison for the 4x4 Square Root topology, by considering injection from G4 and with (right) and without (left) accounting for the inter-express lanes loss-contributions. Every intersection is optimized with standard elliptical tapers.	146
8.6	Insertion-Loss calculation of an optical path of a given optical NoC. . . . .	148

# List of Tables

4.1	Layout-aware properties of topologies under test . . . . .	56
4.2	Parameters used in this work . . . . .	59
4.3	Layout-aware properties of topologies under test . . . . .	68
4.4	Parameters of the simulated architecture . . . . .	80
5.1	Static Power and Dynamic Energy of Electronic and Optical Devices. . . . .	89
5.2	Messages generated by the coherence protocol. . . . .	92
6.1	Photonic Components and Device Parameters. . . . .	107
6.2	Static and Dynamic Power of Electronic Devices. . . . .	109
6.3	Parameters of the simulated architecture. . . . .	111



# Thesis Abstract

Many crossbenchmarking results reported in the open literature raise optimistic expectations on the use of optical networks-on-chip (ONoCs) for high-performance and low-power on-chip communications in future Many-core Systems. However, these works ultimately fail to make a compelling case for the viability of silicon-nanophotonic technology for two fundamental reasons:

- (1) **Lack of aggressive electrical baselines (ENoCs).**
- (2) **Inaccuracy in physical- and architecture-layer analysis of the ONoC.**

This thesis aims at providing the guidelines and minimum requirements so that nanophotonic emerging technology may become of practical relevance. The key enabler for this study is a cross-layer design methodology of the optical transport medium, ranging from the consideration of the **predictability gap** between ONoC logic schemes and their physical implementations, up to **architecture-level design issues** such as the network interface and its co-design requirements with the memory hierarchy.

In order to increase the practical relevance of the study, we consider a consolidated electrical NoC counterpart with an optimized architecture from a performance and power viewpoint. The quality metrics of this latter are derived from synthesis and place&route on an industrial 40nm low-power technology library. **Building on this methodology**, we are able to provide a realistic energy efficiency comparison between ONoC and ENoC both at the level of the system interconnect and of the system as a whole, pointing out the sensitivity of the results to the maturity of the underlying silicon nanophotonic technology, and at the same time paving the way towards compelling cases for the viability of such technology in next generation many-cores systems.



# Thesis Introduction

Optics could solve many physical problems of on-chip interconnect fabrics, including precise clock distribution, system synchronization, bandwidth and density of long interconnections, and reduction of power dissipation. It may allow continued scaling of existing architectures and enable novel highly interconnected or high-bandwidth architectures.

However, despite the arguments in favor of optics for interconnects, and the promising monolithic integration routes with silicon, there is essentially no practical use today. The high cost targets for introducing this emerging technology, and the low-maturity of basic optical components, do not fully justify this scenario, since they only urge for more compelling cases where the benefits of chip-level nanophotonic interconnection networks can justify the cost barrier removal and a larger investment in technology development. It follows from this that such compelling cases cannot be directly derived from the benchmarking frameworks between optical NoCs (ONoCs) and their electronic counterparts (ENoCs) reported so far in the open literature.

While making an excellent point for the new interconnect technology, they lack of enough practical relevance to push it beyond the boundaries of an elegant research concept. In practice, they tend to deliver overly optimistic results for ONoCs for one or more of the following reasons.

**First**, logical topologies are not well specified, hence preventing in-depth architecture review.

**Second**, the baseline electronic NoCs exhibit naive or unoptimized architectures, hence overlooking that performance or power optimizations of ENoCs are many times far more practical than adopting an emerging technology.

**Third**, specific instances of device parameters are currently meaningless for a fast-developing technology such as ONoCs.

**Fourth**, the fixed power overhead is typically underestimated, directly or indirectly by assessing ONoC designs under high utilization regimes.

**Fifth**, complex designs increase risk in terms of reliability, fabrication cost, and packaging issues. In addition to that, we would like to stress that place&route constraints are typically overlooked in ONoC topology design (hence underestimating layout-induced waveguide crossings and static power), and that electronic network interfaces for ONoC injection/ejection are often not considered in planning resource budgets.

This Thesis aims at a higher level of practical relevance in assessing the potentials of ONoCs for future multi- and many-core systems. This is fundamentally achieved in two ways. On one hand, *we make use of an aggressive electrical baseline*. We consider a realistic design point for the ENoC architecture in terms of complexity and power. Moreover, real synthesis runs of the target ENoC on a 40nm industrial low-power technology will provide the reference quality metrics the competing optical NoC solutions are contrasted with. On the other hand, *the ONoC is designed and accurately characterized based on both accurate physical-layer and architecture-layer analysis*. After pursuing a wide topology exploration based on realistic design points (e.g global connectivity and network partitioning) a wavelength-routed optical Ring topology, whose simplicity can reduce the adoption risk of an emerging technology, is selected to be the perfect antagonist to the electronic NoC. At the physical layer, the increased accuracy in ONoC modeling is achieved by drawing the Ring layout, especially its injection and ejection interfaces. At the architecture layer, the design of the network interface architectures needed to inject/eject electronic packets into/from the ONoC is made, thus capturing typically overlooked sources of performance and power overhead, such as flow control, clock resynchronization, or suitable FIFO sizing.

Another feature of this Thesis is that it *carefully considers fixed-power overheads*, which are a significant percentage of total ONoC power. Static power is especially important in those application domains where the network does not undergo high utilization, but it has to serve sporadic traffic peaks. This is the case of shared memory multiprocessors with distributed last-level cache, implementing hardware support for cache coherence. The use of an ONoC makes sense in this domain only if it can significantly cut down on the to-

tal application execution time, thus burning less static power. This Thesis considers the case study of a directory-based implementation of the MOESI protocol, and derives the requirements for both ENoC and ONoC design. Therefore, the compared interconnect solutions are fine tuned for the kind of messages they are supposed to route.

The enhanced level of accuracy pursued by this Thesis in crossbenchmarking optical vs. electronic interconnect technology primarily aims at *providing guidelines of practical relevance to materialize the nanophotonic concept into an affordable technological solution for the next generation multi- and many-core systems.*



# Chapter 1

## On-Chip Optical Communications

### 1 Penetration of optical links into communications

In order to understand how the penetration of optical links into communications will be in the near future, it is useful to examine its history, and in particular over the last 30 years. Figure 1.1 shows an approximate perspective of the rate of penetration of optics versus the link distance and the bandwidth. The lower horizontal axis represents the first commercial introduction of the optical link. The vertical axis represents the minimum range of the link, and the upper horizontal axis represents the bandwidth per fiber connector. The approximate trend suggests that over the last three decades, optical links have achieved an order-of-magnitude deeper penetration into the interconnection hierarchy (from cross-country trans-oceanic applications down to package /chip systems) every five years. Also, it is worth observing that the penetration rate is impressive, around 100 Gbit/s bandwidth improvement per meter of reduction of the connection. Continuing this trend suggests that during the next years, optical links not only can be expected to reach right to the chip-scale package on a printed circuit board, but also to go into smaller and smaller scale systems (e.g. multi-chip systems based on a silicon on interposer technology, systems-on-chip), **thus transforming the concept from**

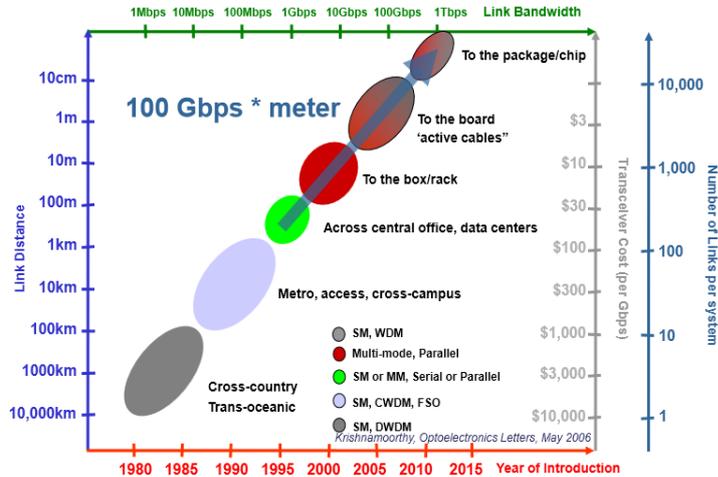


Figure 1.1: Penetration of optical links into communications

### penetration to integration of optical links into systems.

Although there are currently a lot of investments that make optical communication and switching more pervasive in data centers and multi-chip systems based on a silicon on interposer technology, actually there is still an unanswered question: **Will optical links be able to penetrate deeper into smaller systems such as Systems-on-Chip?**

In this context, for sure, optical links have to contrast the current on-chip electrical wires which are by definition inexpensive and hard to beat within certain link distances. However, we should not forget that long on-chip electrical interconnects suffer from many physical effects such as crosstalk, reduction of timing skew in signal, impedance matching, which make the optical on-chip link the ideal candidate to overcome such a limitations. In order to answer the previous question, is it possible to learn something from the past? The optical interconnect technology has historically penetrated systems based on the following paradigms:

**At first**, there were superior point-to-point connections which have implied higher cost to replace the traditional electrical wires. However, this was done because electrical links were not able to cope with the impressive bandwidth

**Current on-chip communication actors are NETWORKED!  
(Network-on-chip communication paradigm)**

### Electronic Network-on-Chip

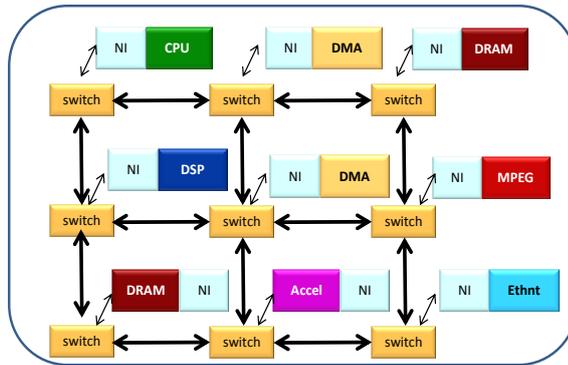


Figure 1.2: Schematic of an Electronic Network-on-Chip (ENoC).

and power requirements. In opaque telecommunication networks, the optical links were used for communication while the electrical counterpart for the switching functionalities (e.g. SFP transceivers).

**Only later**, in transparent telecommunication networks, (e.g. OCS in data-center networks), to improve the quality metrics and remove the higher cost due to the electro/optical and opto/electrical conversions mainly localized at the system boundaries, the optical technology was extended for switching functionalities.

**Presently**, there is a growing interest in optical point-to-point off-chip links and networks (smaller scale systems). Some examples on this topic are: the oracle macrochip [17], hybrid electro-optical controllers for DRAM [30] and also High bandwidth I/O systems [2]. The main driver for such a systems is given by the bandwidth density.

**For deeper chip integration**, will we expect the same story? and how will the feeling change as the off-chip links go over fiber ribbons?

Clearly, the phase transition will be only when the current electronic Networks-on-Chip (depicted in Figure.1.2), built on electrical wires and routers, will not be able to cope with the increased bandwidth demand and power costs

needed for next generation multi-core systems. **In this direction, all of designers and engineers will be searching for compelling cases which make silicon nanophotonic the viable technology for future integrated systems.**

## 2 On-Chip Optical Communication: Why?

Photonic Interconnect Technology is considered a promising way of relieving power and bandwidth restrictions in next generation multi-and many-core integrated systems. Optics could solve many physical problems of on-chip interconnect fabrics, including precise clock distribution, system synchronization (allowing larger synchronous zones, both on-chip and between chips), bandwidth and density of long interconnections, and reduction of power dissipation. Optics may relieve a broad range of design problems, such as **crosstalk, voltage isolation, wave reflection, impedance matching, and pin inductance** [18]. It may allow **continued scaling of existing architectures and enable novel highly interconnected or high-bandwidth architectures.**

Silicon photonics has advanced substantially in recent years and has demonstrated many of the key components for the implementation of future optical networks-on-chip (ONoCs) in an integrated CMOS process [13].

Such components include power-efficient laser sources, low-loss waveguides, high-bandwidth modulators, broadband photonic switches, and high-sensitivity photodetectors. The improvement of the quality metrics of these components, as well as the integration route with CMOS manufacturing processes, are being relentlessly pursued.

## 3 Silicon Photonics as a Technology Enabler

Silicon is a well-known material used in microelectronic chips based on (CMOS) Complementary Metal-Oxide-Semiconductor technology. Silicon photonics offers the compatibility with standard CMOS fabrication processes, enabling dense integration with advanced microelectronics. The capability of silicon photonic devices to be integrated into complex platforms, coupled with decades

of high-quality development driven by the microprocessor industry, allows their low-cost and mass-volume production. Silicon photonics provides also an excellent high index contrast between the refractive index of the core (typically 3.5 for crystalline silicon) and the above cladding (typically 1.5 for silicon dioxide). This high index contrast generates higher optical modes confinement, so that the optical signal can be easily guided by devices with sub-wavelength dimensions. Hereafter an overview of all silicon photonic devices of interest for optical NoC implementation are presented, starting from optical links up to laser sources and photodetectors [14].

### 3.1 Optical Links

The optical link is the fundamental building block that must be used to guide the high speed optical signals from the photonic source up to the receiver. The optical link is commonly referred to as waveguide in the optical domain. Recently, sub-micrometer crystalline silicon waveguides [1] have been an excellent option for optical links. Such a structure is able to propagate parallel wavelengths with terabit-per-second data rates throughout the whole chip. Thanks to these appealing properties, it is possible to further build straight, bend, and crossing waveguides, as well as couplers, thus providing all the basic structures for optical communication channels. From experimental characterization, it has been demonstrated that crystalline silicon waveguides are able to deliver data rates up to 1.28 Terabit/s, including 32 wavelengths modulated at 40 Gbit/s each, through a communication link of 5 cm [1]. Waveguide crossings (i.e., intersections of two waveguides) represent the major source of optical power degradation across optical paths although they cannot be really avoided on a single plane chip. Non-negligible attenuations are incurred across optical paths also in terms of propagation loss (in straight waveguides), or bending loss (in bending waveguides). In the open literature, two-dimensional tapers have been proposed in an attempt to minimize the crossing loss across optical paths. Among the most relevant ones, it is worth mentioning the standard elliptical taper [60] and the MMI (Multi-Mode-Interference) taper [61]. In contrast to sub-micrometer crystalline silicon waveguides, deposited silicon nitride waveguides offer many advantages for integrated photonics. Unlike crystalline silicon, the antagonist

silicon nitride can be deposited in multiple layers, similar to electronic wires. The latter case has the capability of eliminating in-plane waveguide crossing losses, once vertical optical couplers are in place [74]. Experimental results show that the transmission of high-speed optical data through a deposited silicon nitride waveguide can achieve 1.28 Terabit/s (as usual including 32 wavelengths modulated at 40Gbit/s each) throughout a 4.3 cm silicon nitride waveguide [2].

## 3.2 Modulators

The silicon electro-optic modulator is an essential device for photonically-enabled on-chip links, since it performs high-speed conversion of an electrical signal into an optical one. It encodes data on a single-wavelength that can be then combined with additional optical signals through wavelength-division-multiplexing on the same physical medium, thus resulting in a cohesive wavelength parallel optical signal. Crystalline silicon microring resonator electro-optic modulators are the most recently used devices among those presented in the open literature. They consist of a microring resonator configured as p-doped-intrinsic-n-doped (PIN) carrier injection device. The standard operation of these devices relies on non-return-to-zero (NRZ), and on on-off-keyed (OOK) modulation signals. To achieve high modulation rates that are typically limited by carrier lifetimes, modulators are driven using a particular mechanism called pre-emphasis method [3]. The electro-optic modulator has also been proposed for polycrystalline silicon [4]. The grain boundaries inherent in the material result in increased optical loss due to scattering and absorption, which end up reducing free-carrier lifetime, and may increase the intrinsic speed of the modulator accordingly. Unlike crystalline silicon, polycrystalline silicon can be also deposited and stacked with other silicon photonic materials for multi-layer integration. Finally, modulators can be also embedded across arrays (silicon electro-optic modulators arrays), so to deliver a major bandwidth boost along the communication channel. Hence, at the output stage of each array, the optical data stream contains multiple wavelengths ready to be transmitted throughout the interconnection network, ending up at the receiver stage.

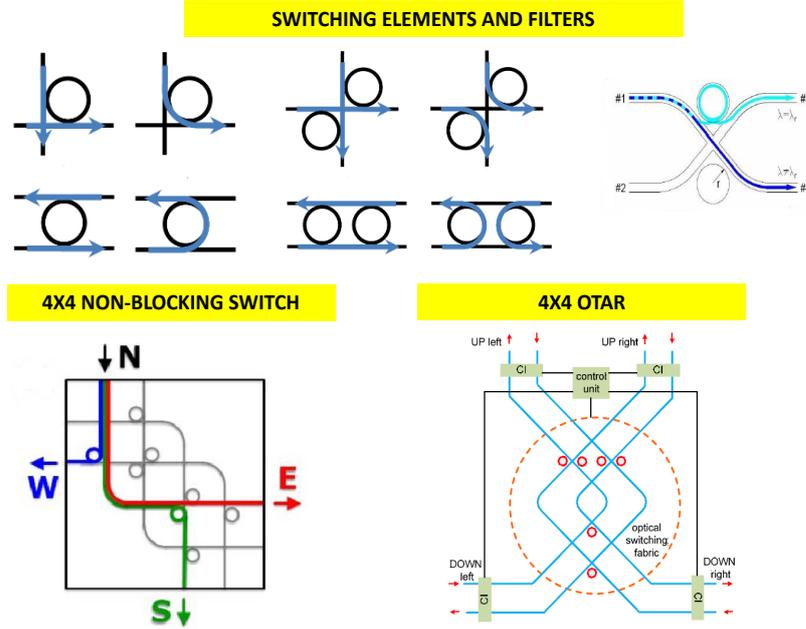


Figure 1.3: Switching Elements and Routers used to build Optical Networks-on-Chip.

### 3.3 Photonic-Switching Elements and Optical Routers for Optical Networks-on-Chip

Broadband PSEs (Photonic-Switching Elements) with 1 or 2 inputs and 2 outputs are the fundamental building blocks of an optical NoC. The former case (one input and two outputs) consists of a microring resonator positioned adjacent to a waveguide intersection. Alternatively, a parallel switching element denoted as 1x2 comb switch has been also presented in the recent literature [15]. Simultaneous switching of 20 continuous-wave wavelength channels with nanosecond transition times has been demonstrated by using the comb-switching technique. 2x2 PSEs instead (two inputs and two outputs) consist of a waveguide intersection and two ring resonators. The switching function is achieved through resonance modulation provided by carrier injection into the micro ring or by designing ring resonators with different radii. The fundamental switching elements introduced above (1x2-PSE, 2x2-PSE) are typically composed to derive higher order switching structures. A 4x4 non-blocking nanophotonic switching node [60] is a clear example thereof.

This optical router may include either 8 1x2 PSEs or a mixture of them, so that each input port is capable of reaching all 3 output ports (because self-communication is not allowed), thus enabling non-blocking functionality. 5x5 Cygnus [12] is another example of strictly non-blocking router for optical NoCs. It consists of a switching fabric, and a control unit which uses electrical signals to configure the switching fabric according to the routing requirement of each packet. The switching fabric is built from the parallel and crossing switching elements. Cygnus uses only 16 microresonators, 6 waveguides and 2 terminators. A 4x4 OTAR (Optical-Turnaround-Router) is an optical router, as always non-blocking, which has been customized for FOnoCs (Optical Fat-tree NoC topologies [59]). It combines a mix of 1x2 and 2x2 PSEs, and is conceived to implement the turnaround routing algorithm typically used by Fat-Tree topologies. All the discussed optical components can be used to build any Optical NoC, such as Mesh, Torus topologies, FOnoCs, as well as multi-stage NoCs and more. Figure.1.3 illustrates some of the aforementioned switching elements and optical routers.

### 3.4 Photodetectors

At the destination front end, a Photodetector is necessary to convert the incoming optical signal into an electrical one. As usual, before sensing the optical signal, a micro-ring-resonator is needed to filter the wavelength-parallel signal, hence treating each component separately. Recently, developments in integrating Germanium Photodetectors with crystalline silicon waveguide have enabled to manufacture many high-performance and CMOS-compatible devices [5, 6], aiming at high-bandwidth (40GHz), high-responsivity (1 A/W), quantum efficiency above 90%, low capacitance (around 2fF), and finally a dark current below 200nA. Another emerging methodology used in the design of photodetectors consists of adopting silicon with crystal defects as the absorbing material. The latest efforts in this field have yielded silicon photodetectors with bandwidth and responsivity higher than 35 GHz and 10 A/W respectively [7]. Similar to modulators, photodetectors can be structured into photodetector arrays. This strategy is very common when a parallel data stream comprised by multiple wavelengths has to be received at the destination stage of a given optical Network-on-Chip architecture.

### 3.5 Laser Sources

For on-chip application, Laser sources can be implemented either on-chip or off-chip, depending on the power and bandwidth requirements of the system at hand, and their trade-offs. Recent emerging technologies continue to mature, and high-quality on-chip lasers compatible with CMOS processes start to appear. Other solutions have been yielded more recently, such as Electrically pumped hybrid silicon lasers and electrically pumped rare-earth-ion lasers on silicon [8]. Alternative solutions leverage on III-V compound semiconductors to produce off-chip laser sources where the light is emitted by the external source and then brought on-chip using couplers. For instance, quantum dot lasers, based on III-V compound semiconductors are typically used in WDM (Wavelength-Division-Multiplexing) applications since they are able to deliver many narrow-spectrum peaks across the frequency range of interest. Opportunely coupled with quantum dot semiconductor amplifiers, these lasers are able to provide several wavelengths within a low RIN (Relative-Intensity-Noise), so that light will be modulated, transmitted and received with error free performance.

### 3.6 3D-Stacked integrated systems (3D-ICs)

Finally, it is worth observing that 3D-Stacked integrated systems (3D-ICs) represent the most likely target for the exploitation of optical interconnect technology. The key reason is that it is a cost-effective solution for the integration of layers manufactured with different technologies, that this way do not need to be made compatible with one another, except for the obvious alignment and inter-layer communication requirements. Across the same vertically-integrated environment we can accommodate processing, memory, and optical layers, thus resulting in a successful strategy to provide low latency, high bandwidth, and cross-layer communications [9, 10], in next generation high-performance multi- and many-core systems.

## 4 Conclusion

However, despite the arguments in favor of optics for silicon chip communications and interconnects, and the success of technology platforms fostering a fables silicon photonics ecosystem [19], ONoCs are fundamentally still at the stage of a promising research concept. At least three reasons can be identified.

**First**, the adoption cost of this technology is still very high, far away from that of the inexpensive on-chip electronic interconnects. This implies that the new interconnect technology will become practically viable only when it will be proven to deliver out-of-reach performance or power figures in the context of compelling use cases. **Second**, technology maturity is currently lagging behind actual industrial standards (e.g., due to thermal sensitivity concerns), and again only compelling cases for silicon nanophotonic links can foster a larger investment on technology development. **Finally**, the availability of mature optical components is not currently supported by mature cross-layer design methods and tools for system design. System designers should be equipped with the needed methodologies and toolflows to do design with the new interconnect technology.

# Chapter 2

## Optical Networks-on-Chip

### 1 Introduction

All devices presented in chapter.1 are key enablers to materialize Optical Networks-on-Chip consisting of multiple optical routers (broadband active switches or passive filters, depending on the routing methodology), that are properly interconnected with each other using silicon waveguides, that in turn may be straight, bent, or crossed depending on the topology requirements. Ultimately, all devices necessary to build an entire on-chip optical communication infrastructure are viable for integration on a silicon chip, thus paving the way for the assessment of the Optical NoC paradigm.

This chapter offers an overview of the most important optical NoC topologies proposed so far in the open literature.

### 2 Optical Networks-on-Chip

The use of optical networks is currently being investigated as a potential methodology to interconnect on- and off-chip components. Currently, optical NoCs are identified as either using wavelength-selective routing or ones that leverage space routing. The main difference lies in the method used to establish the source-to-destination path in the optical medium.

By adopting **wavelength-selective routing**, the switching functionality is implemented with wavelength filters that are arranged throughout the whole network. The optical filters are tuned to allow that each source is routed to

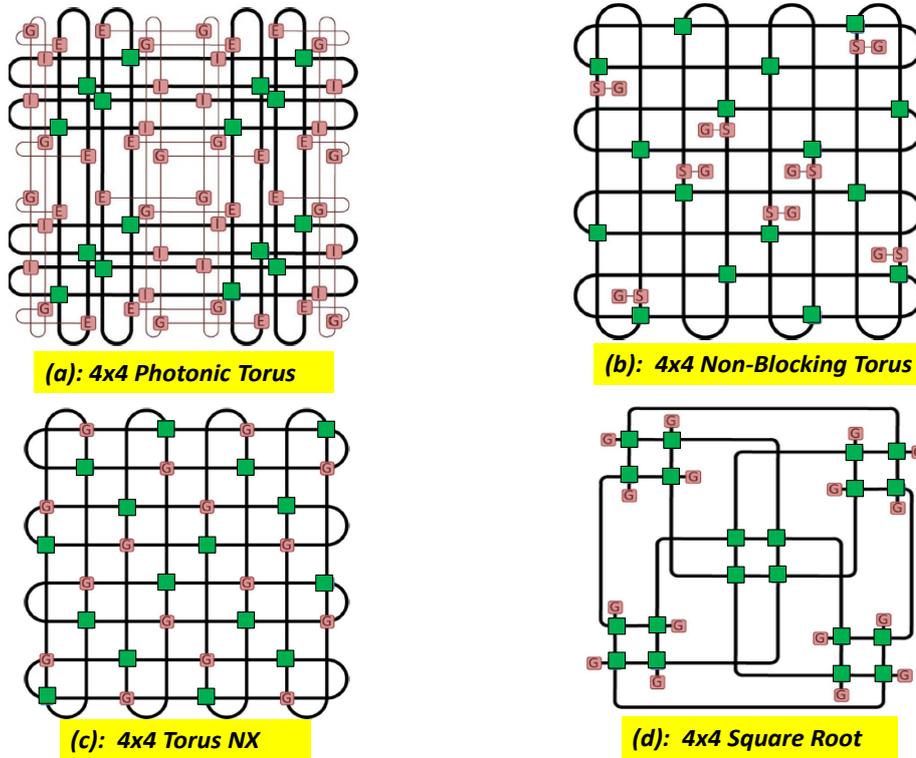


Figure 2.1: Milestones regarding the Space Routed Optical Networks-on-Chip.

each destination through the selection of specific wavelengths. This approach can be described as source routing since the selection of the wavelength at the transmission node determines the entire network path that is used to reach the proper destination node. This routing method enables low-latency communications leveraging on the contention-free property on which such networks are based. However, due to the need of an increased amount of physical resources such routing method is not able to leverage the full throughput that optics can provide.

In contrast, the **Space routing method** concentrates on the use of multi-wavelength transmission to enable messages with high-aggregate bandwidth. These networks are designed to use actively controlled broadband switches to route the whole spectrum of wavelength channels from source to destination. Such optical NoCs rely on a 3D stacking where an electronic control plane, mirroring the photonic network layout, is positioned at the bottom of the optical plane (which lies in the top layer). In particular, the elec-

tronic layer, which is typically realized by an electronic Network-on-Chip, is used to control each broadband switch, and by using a circuit-switching protocol the path setup reservation is accomplished. Once the path is configured, a spatially routed optical network can fully exploit the whole optical spectrum (using WDM technique) to create extremely high-throughput links. However, the required circuit-switching protocol provides an overhead that creates longer latencies with respect to the wavelength selective routing methodology.

### 3 Space-Routed Optical Networks-on-Chip

Figure.2.1 and Figure.2.2 show some of the most famous space-routed optical NoCs (SP-ONoCs) recently proposed in the literature. The following sections will present the main features of such optical NoCs.

#### 3.1 4x4 Torus

Two previously proposed topologies are the Torus and a Non-Blocking Torus, as shown in Figure2.1 (a) and (b), respectively. We define a node (green box) as the logical switching point on the network, while an access point (identified as G) is a gateway. The latter one represents the network user (e.g., a processor element) that can start or receive a certain transmission. The nodes are implemented with the nonblocking 4x4 optical switch as described in chapter.1. The primary folded-torus path in both networks is illustrated with thick lines to represent two waveguides forming a bidirectional link. The remaining thinner lines and blocks ( I,E and S) indicate the location of additional waveguides and switches that compose the access network, which is needed to enter and exit the tori.

The main difference between the two topologies is the way in which access points are mapped to nodes. The Figure2.1 (a) shows the 4x4 Photonic Torus with 16 access points. Switching and access points are green and G boxes respectively. I and E instead represent the switches used to inject and eject messages into and out of the network respectively.

Figure2.1 (b) shows the 4x4 Non-Blocking Torus with eight access points.

S labels indicate combined injection-ejection switching points. Torus has an access point mapped to every node, while the Non-Blocking Torus is limited to two access points on each row and column of nodes in the Torus in order to achieve a strictly nonblocking functionality. For example, an 8x8 Torus would allow 64 access points in a normal configuration but would only allow 16 access points in a nonblocking configuration.

### 3.2 4x4 Torus NX

The Torus NX topology (see Figure 2.1 (c)) is designed to preserve the connectivity and scalability of the original Torus topologies with the main advantage of minimizing the overall insertion loss. In contrast with the original Torus, which required a complex access network to facilitate injection and ejection from the network, Torus NX uses an optimized gateway design, which splits the access point into two blocks for modulation and detection, and circumvents adding any additional crossings to the Torus through the use of the parallel 1x2 PSE. The modulation block enables a message to be injected north or south, while the detection block can collect signals coming from the east or west direction. This scheme is well suited for dimension-ordered routing which is the assumed routing for this topology.

### 3.3 4x4 Square Root

The Square Root topology is also designed with fewer waveguide crossings and switches in mind by simplifying the entire network into only using 4x4 nonblocking switches. In addition to the axioms used to reduce insertion loss in the physical layer, the Square Root leverages on an hierarchical organization to simplify routing and path multiplicity between units to increase performance. The Square Root is constructed recursively beginning with a 2x2 quad, which does not feature waveguide crossings outside the 4x4 switches. A 4x4 Square Root is composed of four sets of quads, and is shown in Figure 2.1 (d) , connecting quads through central switches and interquad connections. Similarly, an 8x8 Square Root can be constructed from four 4x4 Square Roots. This recursive construction can be used to build any size square topology with dimensions equal to any positive integer power of two.

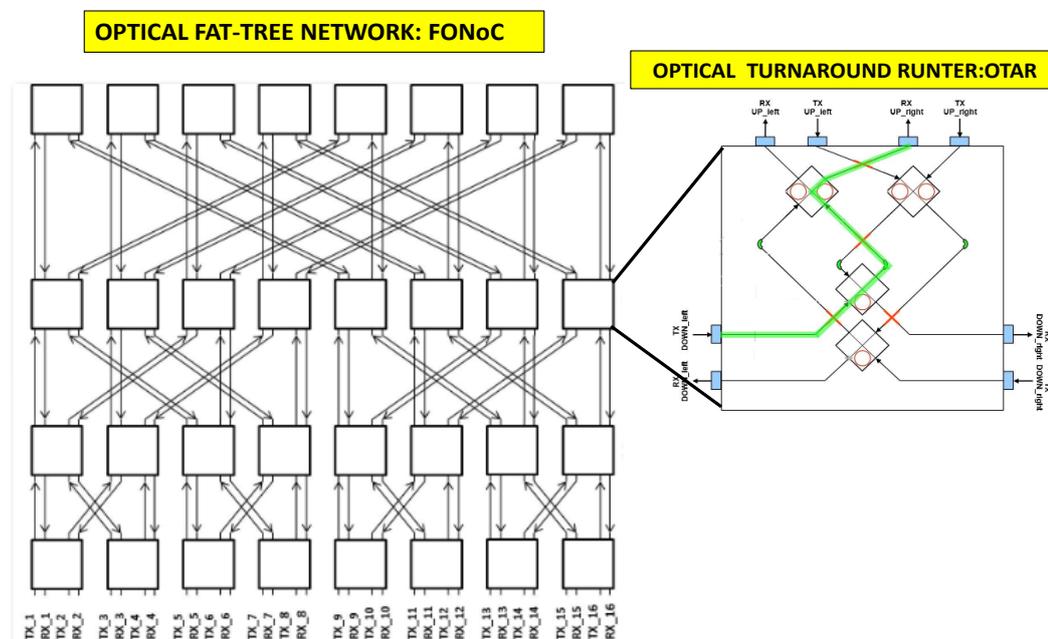


Figure 2.2: The Optical Fat-Tree: FONoC.

### 3.4 The Optical Fat-Tree: FONoC

Figure 2.2 shows the optical Fat-Tree topology which is named FONoC. Different from other optical NoCs, FONoC transmits both payload and control packets on the same optical network. This leads to a lower cost for building a separate electronic NoC for control packets. The hierarchical network topology of FONoC makes it possible to connect the FONoCs of multiple MPSoCs (Multi-processor Systems-on-Chip) and other chips, such as off-chip memories, into an inter-chip optical network, and form a more efficient multiprocessor system.

As shown in Figure 2.2 FONoC is based on Fat-Tree to connect OTARs (previously described in chapter 1) and processor cores. It is a non-blocking network, and provides path diversity to improve performance. Processors are connected to OTARs by using optical-electronic and electronic-optical interfaces (OE-EO), which are needed to convert signals between optical and electronic domains. **FONoC(n,m)** connects **m** processors using an **n-level**

Fat Tree. There are  $m$  processors at level 0 and  $m/2$  OTARs at other levels. To connect  $m$  processors, the number of network levels required is  $n = \log_2(m+1)$ . While connecting with other MPSoCs and off-chip memories, OTARs at the topmost level route the packets from FONoC to an inter-chip optical network. In this case, the number of OTARs required is  $m/2 \log_2(m)$ . If an inter-chip optical network is not used, OTARs at the topmost level can be omitted. In this case, only  $m/2 \log_2(m-1)$  OTARs are needed. As can be seen, in 2.2 each optical interconnect is bidirectional, and includes two optical waveguides.

## 4 Wavelength-Routed Optical Networks-on-Chip

Figure 2.3 shows the most appealing Wavelength-Routed Optical Networks-on-Chip recently suggested in the literature. A different approach is taken by the 4x4 lambda router [41], the milestone switching fabric for Wavelength-Routed Optical NoC topologies. Here, the network routing function is statically determined based on the wavelength of the optical signals. For a given initiator, signal modulated on different wavelengths will be routed differently in the network, and will reach different destinations. Topologies are designed in such a way that signals with the same wavelength originating from different initiators will never interfere with each other.

The appealing property of these topologies is that they enable contention free communication, hence there is no path setup nor contention resolution phase prior to optical packet transmission. This is achieved at the cost of penalizing the bandwidth of each communication stream, although a limited amount of wavelength parallelism is still feasible [38]. Alternatively, spatial division multiplexing can be used. In the lambda router case, with 6 2x2-optical-filters tuned on 4 different wavelengths it is possible to realize 4 filtering stages, resulting in a 4x4 Multi-Stage-Optical Network. Hence, increasing the total number of wavelengths (and in turn the corresponding number of stages), and replicating the number of the 2x2 optical filters, it is possible to derive topologies of arbitrary size. Other switching structures have been proposed

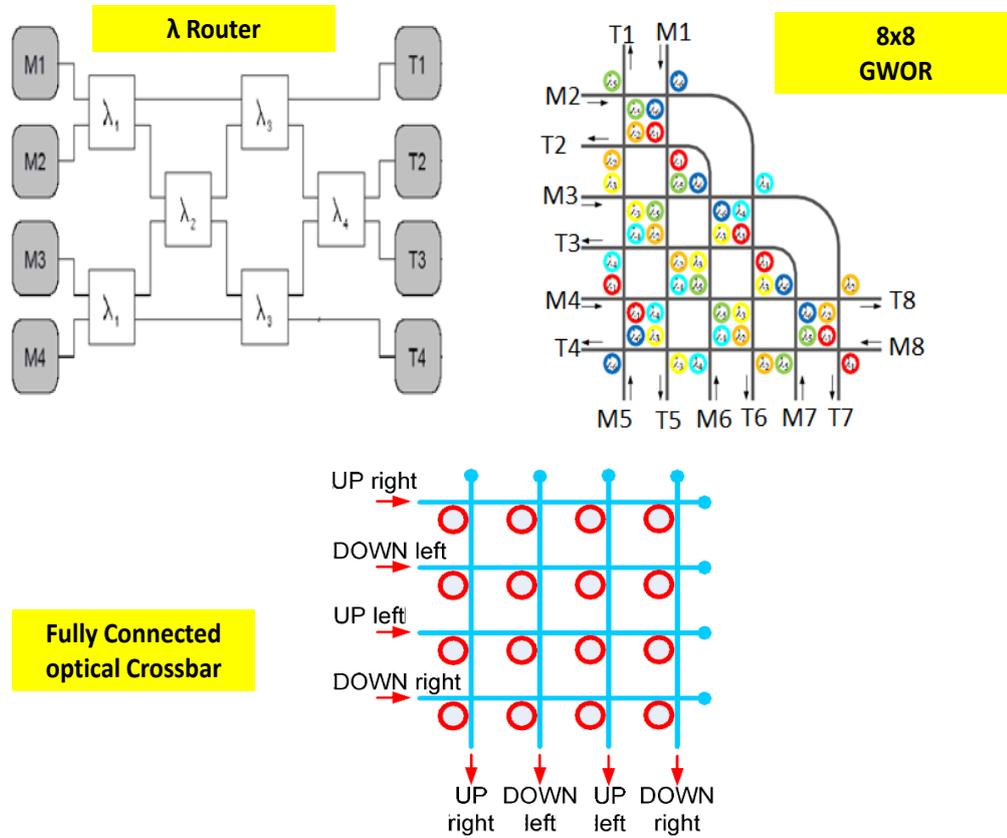


Figure 2.3: Milestones regarding the Wavelength Routed Optical Networks-on-Chip.

that follow the wavelength routing paradigm, such as the 8x8 GWOR (Generalized Wavelength-Routed Optical Router) [28]. This optical structure is capable of enabling 56 contention free optical paths (7 from each input port) thanks to its 48 2x2-optical-filters and 7 wavelengths. It is more suitable to connect initiators and targets distributed across the four cardinal points. The switching fabric of an optical NoC can be also implemented by the traditional fully-connected crossbar. In general, a  $n \times n$  crossbar is composed of  $n^2$  micro resonators and  $2(n - 1)$  crossing waveguides on the critical path. Hence, a 4x4 fully-connected crossbar has four input ports and four output ports and accommodates 16 micro resonators and 7 crossing waveguides on the critical path. An optical Ring network is proposed in [66] by upgrading it to the Spidergon topology for all-optical wavelength routing. The scalability limitations have been overcome in [65], where a two dimensional hierarchical expansion

of the Ring topology is developed. Le Beux et al [56] propose another optical Ring design. By exploiting both techniques spatial and wavelength division multiplexing it is possible to reuse the number of wavelength channels among different physical waveguides, optimize losses, and finally reduce the overall laser power.

## 5 Conclusion

This chapter has provided an overview of the most famous optical NoC topologies both Wavelength-Routed and Space-Routed ones recently proposed in the open literature. In particular, the Wavelength Routed NoCs will be further detailed over the following chapters of the thesis.

# Chapter 3

## Towards Trustworthy Crossbenchmarking Framework between ONoC and ENoC: The Golden Rules

### 1 Pathfinding Requirements

As depicted in figure 3.1 there are three fundamental groups of researchers that are currently involved in the pathfinding effort from the elegant optical NoC concept to an actual technology of practical relevance. The first group is focused on the characterization and optimization of **Silicon Photonic Devices**, and on their monolithic integration with mainstream CMOS manufacturing processes. A relevant gap separates baseline silicon photonic devices with **On-chip Communication Architectures**, that is the topic of the second group of researches, which combine such devices together to materialize higher-order switching structures, complete communication channels, network topologies, routing and flow control methodologies, layout constraints aware physical designs. The last group of researchers is instead involved in the **redesign of an entire system** to take the maximum advantage of the new interconnect technology. Complex and scalable optical interconnects such as Torus, Square Roots [27], hierarchical Wavelength-Routed Optical

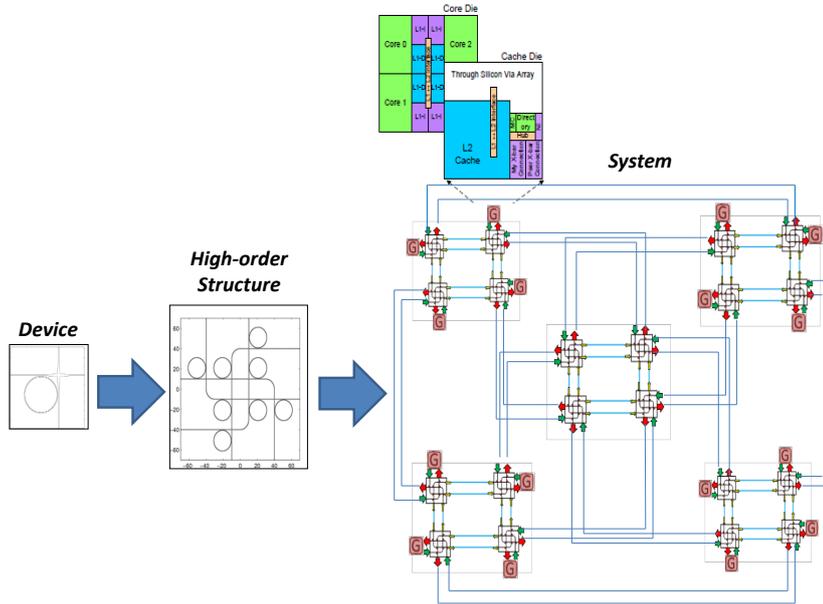


Figure 3.1: The Pathfinding Requirement

Ring architectures [56], as well as Corona and Firefly frameworks [32], [33] have been recently reported in this context. Although many valuable research works have been reported in these emerging research fields, very few structured and coherent methodologies have been proposed so far ([34] is a nice example) to bridge the gap between the above abstraction layers for designing ONoC architectures. Such a pathfinding effort should address two relevant gaps. The first one exists between silicon photonic devices and on-chip communication architectures, which could be referred to as the physical gap. A design methodology addressing this gap should for instance deal with the deviation of physical topologies with respect to their logic schemes, take placement and routing constraints into close account for topology assessment and selection. The second gap separates on-chip communication architectures with system level design frameworks, and could be referred to as the **systemability gap**. The focus here is on the co-design of the optical NoC architecture with the requirements dictated by the target system, and on future generations of such systems.

## 2 An Overview of the Golded Rules

Based on the above considerations, it worth noting that physics and designers should define rules and/or design methodologies in order to be able to bridge both gaps. This chapter in fact goes through a preliminary set of rules, referred to as “**the Golden Rules**”, that should always be followed to effectively design an optical Network-on-Chip. Hereafter are summarized the Golded Rules.

**RULE#1:** Specify a logical topology depending on connectivity requirements requested by the system under test.

**RULE#2:** Account for the Place&Route constraints of a given system.

**RULE#3:** Explore the space of mapping options to nanophotonic devices.

**RULE#4:** Perform an accurate design of network interfaces architectures.

**RULE#5:** Consider an aggressive electronic baseline.

**RULE#6:** Assume a broad range of device parameters.

**RULE#7:** Carefully consider static power overhead.

**RULE#8:** Keep the optical NoC simple to minimize risk.

A brief discussion of them is reported below:

**RULE#1:** In general, any designer should select the type of **ONoC logical topology based on the specific application requested by the system under test**. For instance, the designer could opt for Wavelength-Routed ONoCs (WR-ONoCs) as previously discussed in chapter.2, whether the target system requires low-latency and contention-free full connectivity, where no path-setup reservation is needed. Space-Route ONoCs (SP-ONoCs) will be instead selected for those kinds of applications that require a large transmission bandwidth between cores. As explained in chapter.2, SP-ONoCs need of the path-setup reservation which is commonly accomplished by a dual electronic NoC compared with the previous WR-ONoCs.

**RULE#2:** 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. 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, which raises the design predictability concern for ONoCs as well. Insertion loss, crosstalk and power analysis are important steps to tackle such a concern, 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.

Moreover, as reported in chapter.1, 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, that might break the regularity assumptions of ONoC connectivity patterns, or the floorplanning assumptions they are tied to. Ultimately, the impact of **place&route constraints** might be especially severe for wavelength-routed ONoC topologies.

**RULE#3:** Once selected the set of logic topologies that satisfy the specific requirements of the system under test, it will be then possible to assess such a topologies based on a technology-aware description. In particular, for each topology it will be possible to annotate the overall number of micro-ring resonators (MRRs) and take into account the maximum number of crossing waveguides on the critical path, thus providing an accurate **exploration of the space of mapping options to nanophotonic devices**.

Now, by combining these three rules, the logic topologies should be mapped on the optical layer satisfying the specific place&route constraints and in turn transformed into their physical layouts. Afterwards, by leveraging on the physical-layer properties (e.g. insertion loss critical path, total power

viewpoint), the deviation between the logic topologies and their physical layouts (i.e. **the predicatbility gap**) will be quantified.

**RULE#4:** In order to exploit the potentials of the optical medium in future many-core systems, an accurate modeling, design, and simulation of network interfaces architectures should be taken into account. As previously mentioned in rule number 2, any multi-and many-core system-on-chip (MPSoC) accommodates both electronic and optical networks. In particular, at any transmission stage of the electronic network there is need for interfaces. The latter ones are necessary as the electrical signal incoming from a processor element (PE), before being transmitted into the optical network has be converted into an optical one by means of appropriate electro/optical converters. Conversely, at the reception stage, such an optical signal will be back converted into an electrical one (i.e. opto/electrical converters) to be correctly received at the destination node, (another PE or a memory bank).

In light of this, any optical on-chip communication requires for hardware support, not only characterized by the optical devices such as laser sources, modulators and receivers, but also by all of electronic circuitry for buffering, flow control, and resynchronization functionalities.

In particular, all of electronic and optical devices comprising the network interfaces architectures should be carefully designed and optimization techniques should also be searched for.

**RULE#5:** In order to provide a trustworthy crossbenchmarking frameworks between an electronic NoC and the optical counterpart, an aggressive electronic baseline should be considered.

**RULE#6:** Another feature that should be considered in crossbenchmarking framework is the *fixed-power overhead*, which it has a significant percentage of total ONoC power. Static power is especially important in those application domains where the network does not undergo high utilization, but it has to serve sporadic traffic peaks. This is the case of shared memory multiprocessors with distributed last-level cache, implementing hardware support for cache coherence. Consequently, any designer should carefully account for the

static power overhead of each component of both interconnection fabric and network interfaces architectures by using a broad range of device parameters.

**RULE#7:** Last but not least, any designer should implement an optical interconnection network with the following features: simple and with the minimum risk.

### 3 Conclusion

By meeting the presented rules, **first**, it will be possible to build a trustworthy crossbenchmarking framework between an optical NoC and its electrical counterpart. **Afterwards**, it will be possible to identify the searched compelling cases. Although, this chapter has only presented the golden rules and their main motivations, the following chapters will give a detailed exploration of them. In particular, chapter 4 will go through rules number 1, 2 and 3, quantifying the design predictability gap in optical NoC. The chapter 5 will discuss the rule number 4 examining the network interface architecture. Finally, chapter 6 will focus on the crossbenchmarking framework, thus satisfying rules number 5, 6 and 7.

## Chapter 4

# The Design Predictability Gap in Optical Networks-on-Chip Design

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 chapter 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 chapter 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.

# 1 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. In particular, there might be a profound difference between the logic topology and its physical implementation [35], which fundamentally raises the **design predictability gap** for ONoCs. The design predictability gap is more a concern for wavelength-routed ONoC (WRONoC) topologies rather than for space-routed ones. The wavelength parallelism of the latter ones directly matches the bit parallelism of electronic NoCs, and this explains why topologies proposed for space-routed ONoCs are essentially inspired by those of general purpose ENoCs: meshes, tori, fat-trees, spidergon or recursively built topologies. Technology-specific adaptations of such networks have therefore a good matching with the 2D layout surface from the ground up, facilitated by the fact that link length is not a critical parameter for optical links. In contrast WRONoCs are implemented using wavelength filters throughout the network. These 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 [56] to customized multi-stage networks [41, 52, 53], which often make strong and unrealistic assumptions on master and slave placement or total wirelength to achieve compact and efficient implementation.

This chapter 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 abil-

ity to deliver the same communication bandwidth with the minimum power consumption. In summary, the main contributions of this chapter are:

**A.** Due to the lack of automatic CAD tools for designing optical NoCs, a full custom design (manual design) is performed for the place&route of multiple WRONoC topologies, 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 insertion loss critical path and total power consumption are derived for each topology.

**C.** Optical Ring networks will be compared with topologies relying on optical filters (i.e filter-based), thus assessing their actual need in the context of WRONoC domain. The conclusion on this topic will be supported by preliminary scalability results on the same target system.

**D.** In order to increase the level of confidence of this comparative framework, the chapter does 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), global connectivity, network partitioning for wavelength reuse (across all of subnetworks), and slight topology transformations for more flexible and/or efficient place&route (for the optical crossbar and GWOR topologies).

## 2 3D-Target Architecture

The common experimental setting of practical interest to assess WRONoC topologies is a 3D architecture for multicore processors (see Figure.4.1), consisting of an electronic layer and of an optical one stacked on top of it. Similar architectures are already available in the market, e.g. the Tilera family of multi-core processors [42] which currently features arrays ranging from 16 up to 100 cores. The electronic layer that we assume in this study consists of 32 cores that are linked with each other by means of an electronic NoC with a 2D mesh topology. Also, we assume that cores are structured into 4 clusters of 8 cores each. Again, each cluster has a private gateway to access

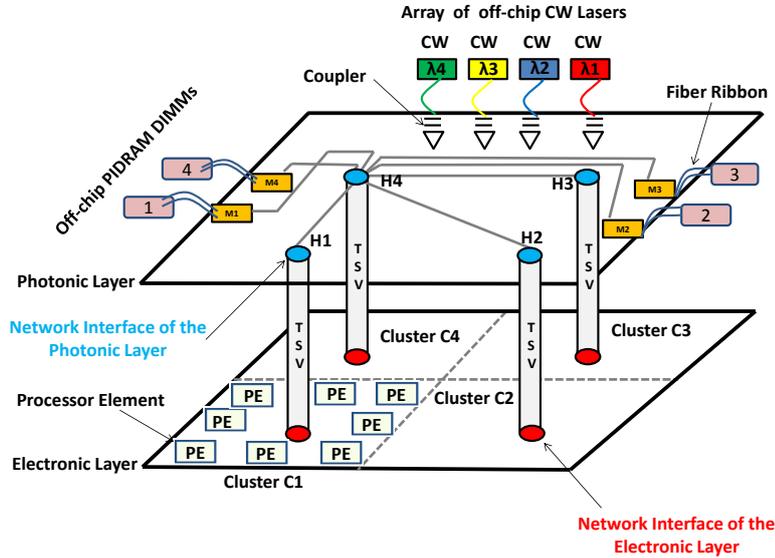


Figure 4.1: 3D-Target Architecture

the above optical layer. We assume an area footprint of  $1.33 \text{ mm}^2$  for each core, and a die size of  $8 \text{ mm} \times 8 \text{ mm}$ .

The optical layer 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 [30]; (c) from a memory controller to a cluster. Also, it 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 4.1).

The optical power is provided by an array of off-chip continuous wave laser sources (CW laser sources) and the multi-wavelength signals are coupled into the chip and brought to the initiators for modulation. As will be proposed later on, the network partitioning allows that the same array of CW lasers can be shared by all the initiators. Therefore, 4 lasers will be needed since every initiator modulates the same 4 wavelengths.

The microarchitecture of memory controllers depends on the specific imple-

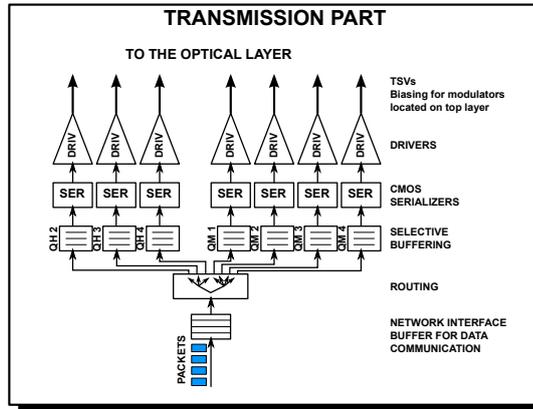


Figure 4.2: Electronic Network Interface: Transmission Side

mentation of the memory sub-system. As an example, in [30] optical command, read and write busses connect the controller to the off-chip photonically integrated DRAM (PIDRAM) DIMMs via a fiber ribbon. In any case, the memory controllers are typically placed all around the chip. Their specific location depends on the position of the PIDRAM DIMMs on the board while at the same time reducing contention (hot spots) in the on-chip interconnect fabric. In this study, we assume 4 memory controllers (M1, M2, M3 and M4) that are located pairwise at the opposite extremes of the chip, as proposed in conventional chip multiprocessor architectures [42], 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. 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 initiators on the same wavelengths.

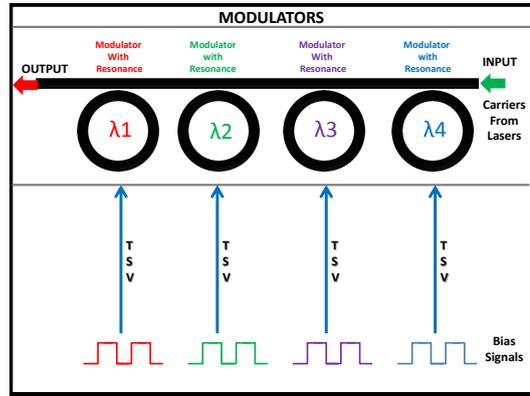


Figure 4.3: Array of modulators in the optical layer

### 3 Electro/Optical & Opto/Electrical Network Interfaces

Let us now focus on the electro-optical network interface (NI), which resides partly in the electronic layer and partly in the optical one. The former section, relative to the transmission part, is illustrated in Figure.4.2. Packets coming from the cluster’s electronic NoC are buffered at the network interface front-end. Based on their destination, they are stored in differentiated buffers (in our target system, there are 7 buffers associated with the other clusters and with the memory controllers). A serializer reads packets from the buffers and feeds them to the drivers. We assume that drivers are directly connected to the through-silicon vias (TSVs) and through them to the modulators on the optical layer. The latest technological developments about 3D-integration enable TSVs with a pitch of  $5\mu\text{m}\times 5\mu\text{m}$  and therefore a large TSV integration density (up to 160K TSVs in a  $10\text{mm}\times 10\text{mm}$  die). Moreover, as reported in [43] [51], the TSVs can deliver high-speed transmission from 1 Gbit/s to 10 Gbit/s. This relevant performance motivates our choice of using them to provide the biasing signal to the optical modulators in the optical plane (see Figure.4.3). The rationale behind this choice is to avoid integrating electronic devices in the optical layer and, therefore, enable low-cost fabrication of this latter, a key requirement to make silicon photonics affordable for the embedded multi-core computing domain. In line with current technology, we assume modulation rates of 10 Gbit/s for each wave-

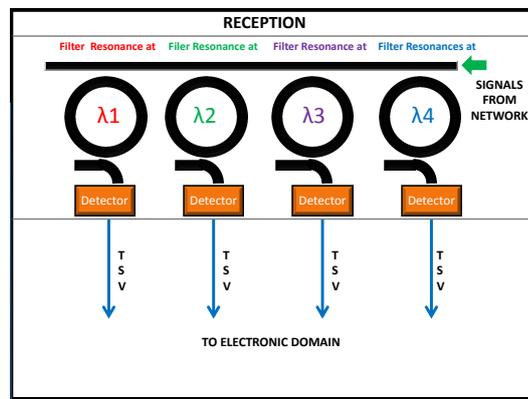


Figure 4.4: Array of filters and detectors in the optical layer

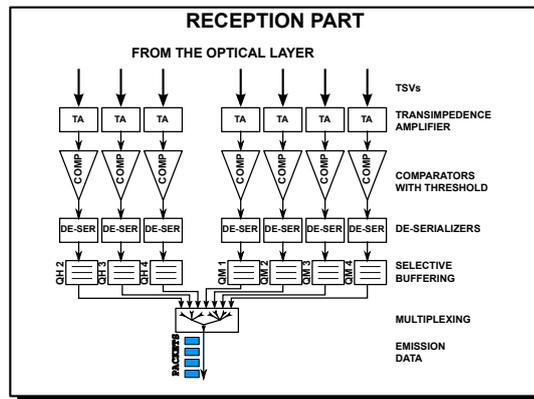


Figure 4.5: Electronic Network Interface: Reception Side

length. Therefore, the injection rate of every hub peaks at 40 Gbit/s. In the optical layer we use WRONoCs, therefore every destination-specific buffer in the electronic NI is associated with a different wavelength in the modulation array. By complementing this with a network made up of add-drop filters, contention-free optical communication is achieved with no latency overhead for arbitration, routing or circuit setup.

The reception part is specular. As depicted in Figure.4.4, the optical layer features an array of add-drop filters for each hub feeding photodiodes that convert the optical signal back into the electrical one. The photodiode's outputs are then conveyed to the transimpedance amplifiers in the electronic layer by means of TSVs. Again, we opt for not placing the electronic devices in the optical layer. As illustrated in Figure.4.5, digital comparators and

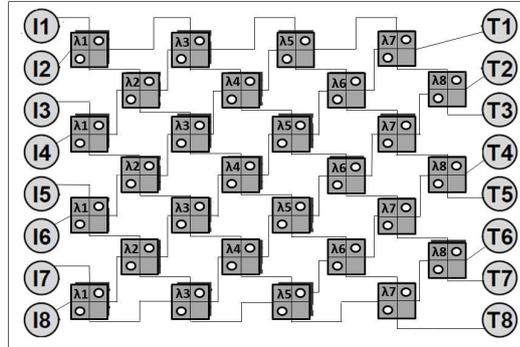
de-serializers complete the domain conversion. Buffers are associated with packet source and from here on the electronic network interfaces come into play (e.g., association of memory responses with memory requests, packetization for the electronic NoC). Although this section only has given an overview of the NI architectures, in chapter.5 they will be further detailed.

## 4 Design Predictability Gap: Logic Scheme vs. Physical Layout

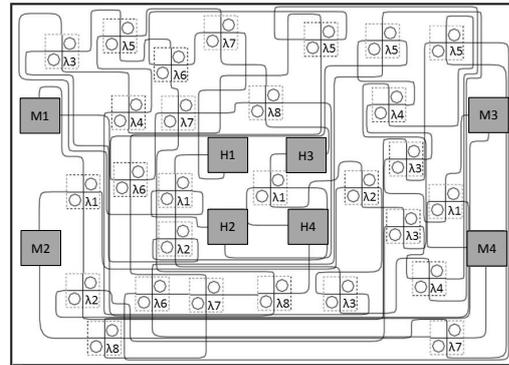
There is still a significant gap between the optical Network-on-Chip concept and a mature interconnect technology with practical relevance, which consequently raises the design predictability gap in ONoCs domain. Hence, insertion loss and total power are important steps to tackle such a concern and to assess the actual feasibility of an optical NoC from a physical layer point of view. Figure 4.6(a) shows the 8x8  $\lambda Router$  logic scheme, (which was proposed by Scandurra et al in [52]), while its corresponding physical implementation is shown in Figure 4.6(b). The latter one is obtained with a manual design [35]. As can be seen, there is a large deviation between the native scheme and its physical layout mainly due to the physical constraints (i.e fixed positioning of hubs and memory controllers). More precisely, the optical path with the major number of waveguide crossings in the physical layout results 9 times larger than the corresponding logic scheme. In fact, it differs from 7 crossings in the logic scheme up to 64 in the physical layout. Finally, it is reasonable to expect that the insertion loss critical path and the laser power requirement too may degrade so much that an appealing logic topology could result not affordable, thus pointing out the **design predictability gap** in optical NoC design.

## 5 Topology Exploration: Global Connectivity

As mentioned in section 2, the proposed optical layer accomodates 8 initiators and 8 targets. The easy way to interconnect them consists of using one 8x8 wavelength-routed optical NoC, thus aiming at global connectivity scenario.



(a)



(b)

Figure 4.6: The predictability gap : (a)  $8 \times 8$   $\lambda Router$  logic scheme, (b)  $8 \times 8$   $\lambda Router$  Real Layout manually generated

Afterwards, the most relevant topology logic schemes that can address this connectivity requirement are discussed, and then implemented in the target system under the effect of layout constraints.

## 5.1 Relative Topology Comparison

Figure 4.6(a) shows the first topology under test: an  $8 \times 8$   $\lambda router$ . In order to interconnect 8 initiators with 8 targets, the network utilizes 8 stages of 4 and 3 add-drop filters (see gray boxes in figure 4.6(a)). The topology reflects the connectivity pattern of unidirectional multi-stage network (MINs) commonly

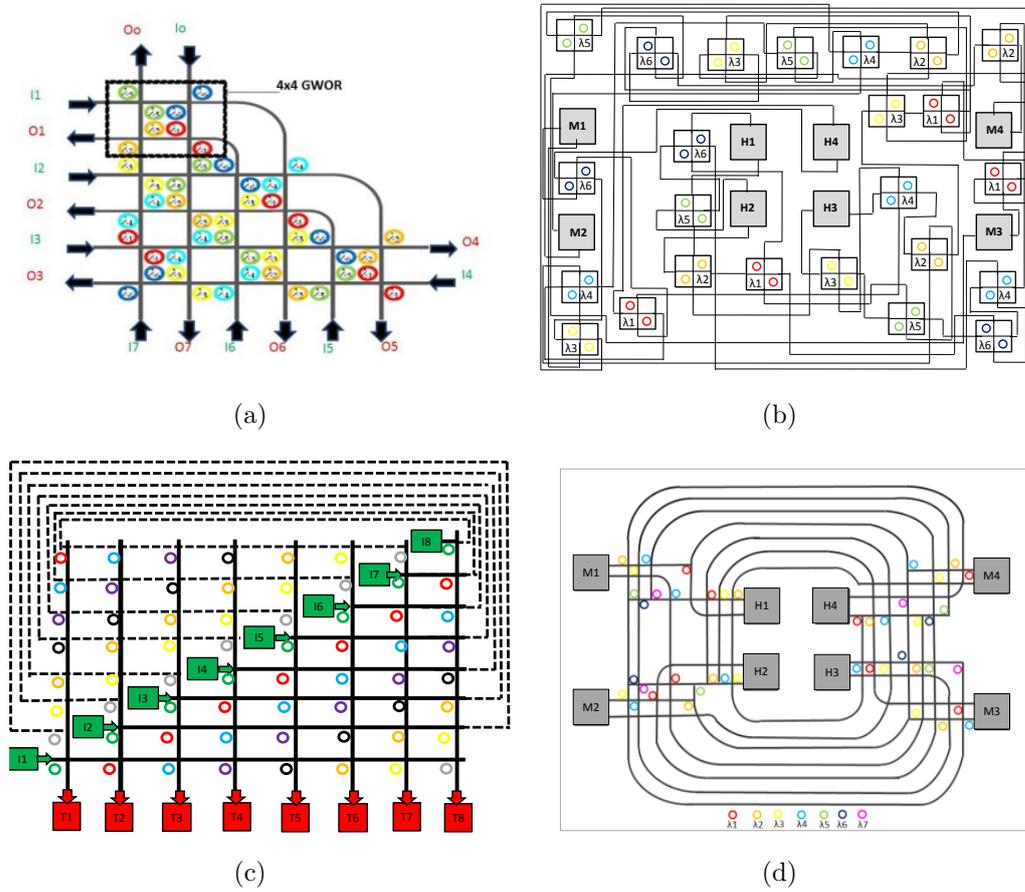


Figure 4.7: (a) 8x8 GWOR logic scheme, (b) 8x8 GWOR physical layout, (c) 8x8 Folded Crossbar logic scheme, (d) 8x8 Folded Crossbar physical layout

used in the electronic domain, with the difference that the inter-stage pattern is closely related to the routing methodology of WRONoCs. Unfortunately, the attractive logic scheme of this topology does not match the actual layout constraints of most real-life systems, where indeed it is almost impossible to find all initiators arranged on one side and all targets on the other side of the chip. At the same time, hubs and memory controllers are both initiators and targets of on-chip communication transactions, hence the physical implementation of this topology implies some degree of folding (see Figure 4.6(b)).

In order to find the best solution for global connectivity, we compare the previous topology with the 8x8 GWOR [28] and an optimized crossbar, here referred to as 8x8 folded crossbar. According to the wavelength assignment

proposed in [28], the 8x8 GWOR logic scheme (see Figure 4.7(a)) is constructed starting from its lower basic cell, the 4x4 GWOR. The latter one is well suited for those cases where initiators and targets are distributed across the cardinal points. In fact, the topology consists of 4 waveguides which intersect each other, with micro ring resonators (MRRs) placed pairwise at each intersection (see colored circles inside the black box of Figure 4.7(a)). Unfortunately, the scaled pattern to an 8x8 network keeps making the same physical placement assumptions, which is not realistic, since it is very unlikely that all cores are placed around a centralized optical interconnect. Moreover, it is worth recalling that, unlike the previous topology, this one is not capable of self-communication.

The physical view of the 8x8 GWOR is illustrated in Figure 4.7(b), and confirms that the placement constraints of the target system are unnatural for the GWOR connectivity pattern, which ends up in a circuitous wiring making the original pattern hardly recognizable. Again, the waveguide crossings arising as an effect of the 2D surface mapping are apparent.

Finally, an 8x8 optical crossbar is considered. This topology places MRRs at each intersection of a point-matrix, thus establishing connections between a given initiator and the desired target. Although considered quite inefficient in abstract analysis frameworks, the topology lends itself to an interesting optimization already in its logic scheme. In the original topology, every initiator delivers optical signals to targets in a given order. By changing this order for every initiator (see Figure 4.7(c)), then we apparently cause a waveguide length overhead. However, this is only an apparent effect of the logic scheme, since the actual layout is in contrast facilitated: every initiator can in fact drive an optical waveguide that enters a Ring-like topology for dispatching of optical packets to the possible destinations (Figure 4.7(d)). Clearly, the handcrafted layout of the folded crossbar is much more regular than the 8x8  $\lambda Router$  and the 8x8 GWOR, and MRRs are clearly positioned close to communication targets for wavelength-selective ejection of optical signals. In previous comparison frameworks, such physical-level details are typically omitted, and considered quality metrics include mainly the number of MRRs and the maximum number of waveguide crossings on the logic scheme. As shown in table. 4.1, at logic scheme level the 8x8 Folded crossbar features

Table 4.1: Layout-aware properties of topologies under test

	Total # of Wavelengths	Path Max # of Crossings L.Scheme	Path Max # of Crossings Layout	Total # of MRRs L.Scheme	Total of# MRRs Layout
8X8 $\lambda$ -Router	8	7	64	56	48
8X8 GWOR	7	10	72	48	36
8x8 Folded Crossbar	7	14	22	64	44

the largest number of MRRs (64), as opposed to 56 for the 8x8  $\lambda$ Router and 48 for the 8x8 GWOR.

It worth noting that, the 8x8  $\lambda$ Router and GWOR accommodate 2x2 optical filters while the 8x8 Folded Crossbar is only based on the 1x2 ones. The ranking is exactly the opposite when the number of waveguide crossings is considered: the  $\lambda$ Router exhibits 7 crossings as opposed to the 10 ones of GWOR and the 14 ones of the optical crossbar.

Unfortunately, this analysis methodology is only partially informative and even misleading. Physical layer and layout analysis are required to assess the actual trade-offs. This way, as self and memory-to-memory communications are not allowed, redundant MRRs can be removed. In particular, 8 MRRs were removed in the 8x8  $\lambda$ Router and the 12 ones in GWOR.

The Folded crossbar benefits of this optimization much more than other solutions (20 MRRs are removed). After applying the MRRs optimization, GWOR continues to show the minimum number of MRRs, (36), followed by the 44 ones of the Folded crossbar, while the 8x8  $\lambda$ Router still shows the maximum number of micro rings (48). Finally, in terms of number of wavelengths, it is very important to keep in mind that optical crossbar and GWOR utilize 7 distinct wavelengths (i.e., 7 Continuous-Wave Laser Sources) by construction to deliver full and global connectivity vs. the 8 ones of an 8x8  $\lambda$ Router. All these layout-aware properties are summarized in table.4.1

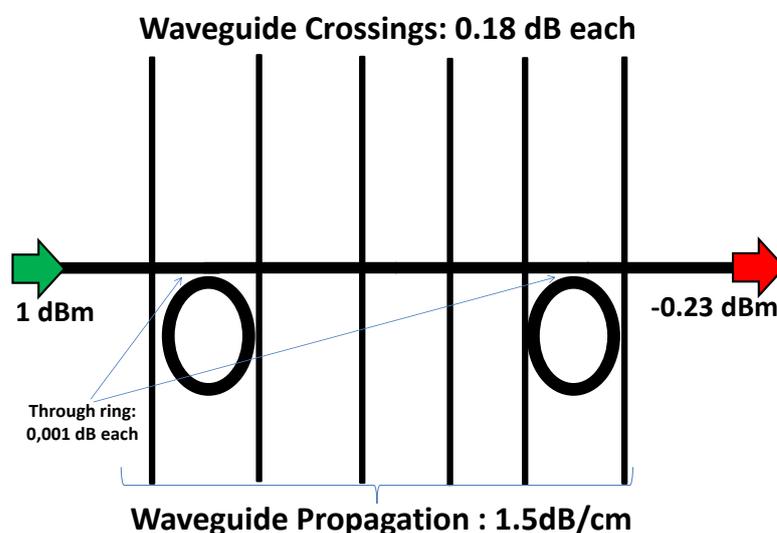


Figure 4.8: Calculation of Insertion loss for a small Network Segment

## 5.2 Physical Layer Analysis

In this section the presented topologies will be compared in terms of insertion loss and power consumption.

### (a) Insertion loss Analysis

Figure 4.8 shows a simple example of a signal injected into a network segment at 1 dBm and being received at -0.23 dBm. The optical signal experiences a propagation distance of 1 mm, passing by 2 micro-ring resonators, and going through 6 waveguide crossings. The overall insertion loss in this case is 1.23 dB. It worth noting that the insertion loss is the most important physical metric that must be quantified to determine the laser power that guarantees a predefined bit error rate at receivers. In fact, once  $IL_{max}$  is obtained (the maximum insertion loss across all wavelengths and optical paths) and the detector sensitivity is known ( $S$ ), it is possible to evaluate the lower limit of optical laser power ( $P$ ) to reliably detect the corresponding photonic signal at the receiver end. We firstly calculate the worst case  $IL_{max}$  across the entire global network, and then we make the practical assumption that such a worst case  $IL_{max}$  dictates the power required by all laser sources. Our study assumes the loss parameters summarized in table.4.2. We rely on a Simulink

simulation framework (see chapter.8) to assess physical metrics of optical NoCs by modeling every single path of a given topology while accounting for the above loss parameters. Finally, we obtain the corresponding insertion loss as a sum of all components encountered in the path under test such as straight, bend, crossing waveguides and drop-into ring loss. The topology models assume die sizes of  $8mm \times 8mm$ .

Figure 4.9 shows insertion loss deviations between logic and physical ONoCs for all topologies considered in this comparison framework. We assume also that the standard elliptical taper[60] is optimized at every waveguide crossing. The insertion loss critical path is more than 6x worse in two physical networks out of three compared with the corresponding logic schemes (see blue and red bars in Figure 4.9). Especially GWOR suffers from 72 waveguide crossings against the 10 expected ones. The  $\lambda Router$  reports 64 crossings vs. 7, thus preserving its superiority over GWOR in relative terms.

Surprisingly, the Folded crossbar maps more efficiently to the target placement constraints, although it is frequently discarded in abstract analysis frameworks, which only consider the logic schemes and the abstract properties. The physical implementation is so efficient (i.e., only very few additional crossings (8) from layout constraints) to offset the inherently higher number of waveguide crossings of the logic scheme.

More in details, propagation loss is a significant contribution in the folded crossbar topology, indicating that the critical path now is both waveguide- and crossing-dominated. Due to its long optical waveguide of 25.5 mm and 22 crossings, the Folded Crossbar achieves a critical path of 15.3 dB. On the contrary, for the 8x8  $\lambda Router$  (33.3 dB) and GWOR (37.5 dB) only crossings have been computed since their contribution is dominant in the breakdown. Obviously, if had we accounted for their propagation losses too, their current huge gap with the crossbar would have become even worse, without providing any significant novelty to the discussed results. Finally, in terms of insertion loss critical path we saw a reverse trend between logic topologies and their physical implementations. In fact, the worst logic scheme topology (F. crossbar) becomes an interesting solution when placed into a real scenario (7.28 dB vs. 15.3 dB). In contrast, the most promising logic scheme topologies ( $\lambda Router$  (3.64 dB), GWOR (5.21 dB)) when placed into

Table 4.2: Parameters used in this work

Parameters	Value	Devices	Features
Propagation-Loss [27]	1.5dB/cm	Laser	CW( Continuous Wave) PLE=20% (Laser efficiency)
Bending-Loss[27]	0.005dB		PCW=90% (Coupling Laser-Link)
Crossing-Loss Optimized by Elliptical Taper[84]	0.52dB	Modulator	Si Disk $\beta = 20\%$ (Launch efficiency) Dyn. Dissipation=3fj/bit Static Power=30 $\mu$ W Vdd=1V
Optimized by MMI Taper[84]	0.18dB		Modulator Power depends on ILmax [63]
Drop-Loss Optimized by Elliptical Taper[84]	0.013dB	Detector	CMOS(45nm) hybrid silicon receiver S=-17dBm, (BER=10 <sup>-12</sup> @10Gbit/s) Power=3.95mW [62]
Optimized by MMI Taper[84]	0.0087dB	Optical Filter	Thermal-Tuning: 20 $\mu$ W/ring [27]

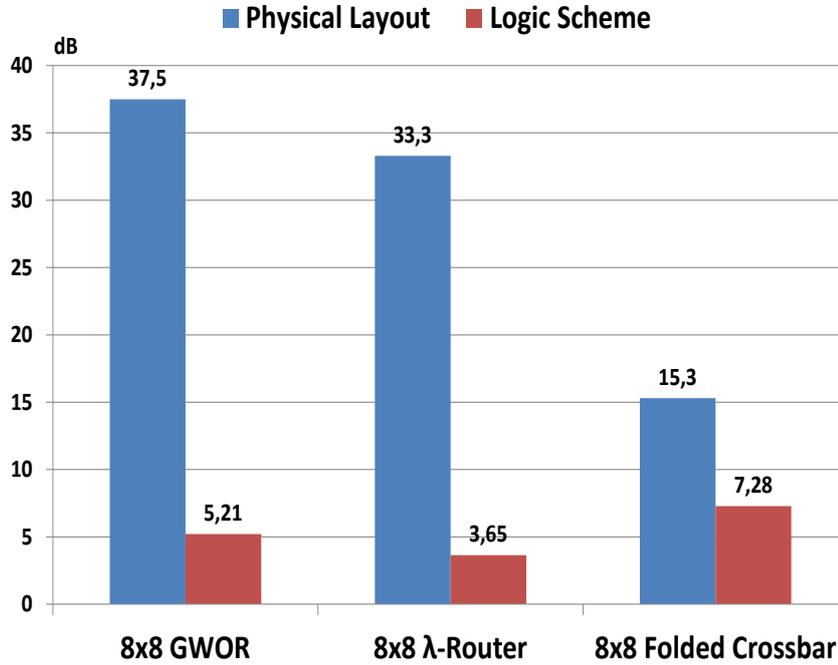


Figure 4.9: ILmax Contrasting: Physical Layout vs. Logic Scheme

their physical layouts are heavily penalized by the layout constraints, thus resulting clearly inefficient with respect to the optimized crossbar.

### (b) Power Analysis

By using the critical path insertion loss, it was then possible to derive the needed laser power to get a bit error rate of  $10^{-12}$  at optical receivers with a sensitivity of -17dBm [62]. By also considering the contribution of modulator and detector power and of thermal tuning that are listed in table.4.2, we derived the total power results consumed across the maximum bandwidth of 440 Gbit/s (see Figure.4.10).

The total power of GWOR is larger than that of other topologies, even if the  $\lambda$ Router features one laser source more than GWOR and the Folded Crossbar to provide the same (full) connectivity. More precisely, the total power of the  $\lambda$ Router topology is 2.47x lower than the GWOR one. The Folded crossbar turns out to be the most power efficient solution. It consumes only

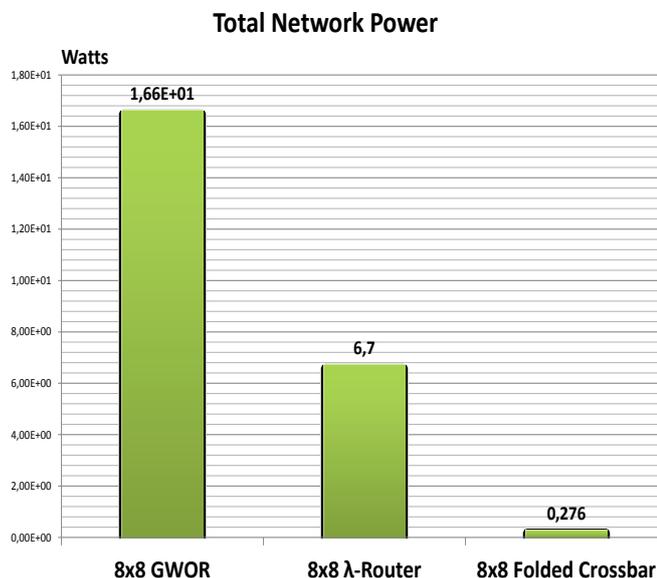


Figure 4.10: Total power Contrasting: Physical Layout vs. Logic Scheme

276 mW, almost 2 orders of magnitude lower than GWOR (16.6 W). Since GWOR and  $\lambda Router$  feature higher insertion loss than Crossbar, their major contributors in the power breakdown are laser sources (28%) and modulators (72%). Moreover, it worth noting that the overall modulator power is higher than laser power since their amount is larger than the laser one (44 vs 7/8). It should be recalled that the modulator power partly depends on the input optical power[63] that in turn depends on  $IL_{max}$ . In contrast, the receiver power mostly contributes in the power breakdown (63%) of the Folded Crossbar due to its lower insertion loss. In the end, these results indicate that GWOR and  $\lambda Router$  are infeasible for the placement constraints of the target 3D-System, while Folded Crossbar turns out to be the best topology.

### 5.3 Comparison with an Optical Ring Topology

The layout of the folded crossbar is very similar to a Ring topology. The latter one is in fact the simplest connectivity solution among all network topologies presented in the open literature, and apparently the less sensitive one to place&route constraints. The only one way to assess whether the Folded

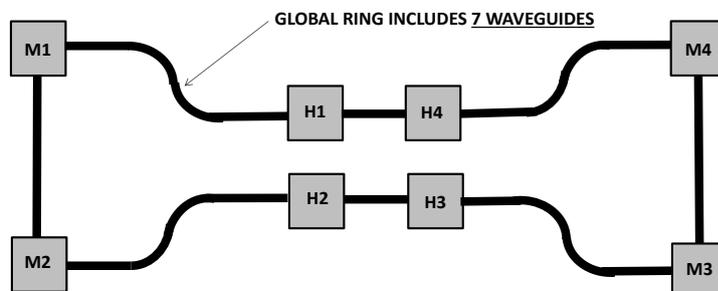


Figure 4.11: Optical Ring Physical Layout

Crossbar is the most efficient solution for the target system consists of comparing it with an actual Ring topology. In order to correctly compare these topologies, a Ring is designed assuming the same number of wavelengths utilized in the crossbar, i.e., 7. The use of multiple ring waveguides (i.e., spatial division multiplexing) and the reuse of wavelengths right across waveguides it is the only way to meet this requirement [56]. Figure.4.11 depicts the real layout of the Ring topology after manual placement of the 3D stacked optical layer under test with the given physical constraints (i.e., hubs in the middle and memory controllers positioned at the opposite extremes). Clearly, this topology better fits the target constraints. Essentially, it works like a bus, in which multiple waveguides are contained into it. In this case, 7 parallel waveguides are needed to deliver full and contention-free communication parallelism. Figure.4.13 illustrates the post-layout insertion loss critical path comparison between the 7-way Ring and the 8x8 Folded Crossbar. The 7-way Ring achieves 7.75 dB insertion loss against 15.3 dB of the Crossbar on the critical path, thus resulting 50% more power efficient. The key reason lies in the fact that the 7-way Ring provides less wiring length on the critical path (2cm vs. 2.55 cm in the Crossbar). Moreover, the crossbar has 22 waveguide crossings (localized in the optical network), while only 9 crossings are there in the 7-way Ring. Even if in the Ring topology there are no intersections in principle, they are actually needed at initiator interfaces to connect to

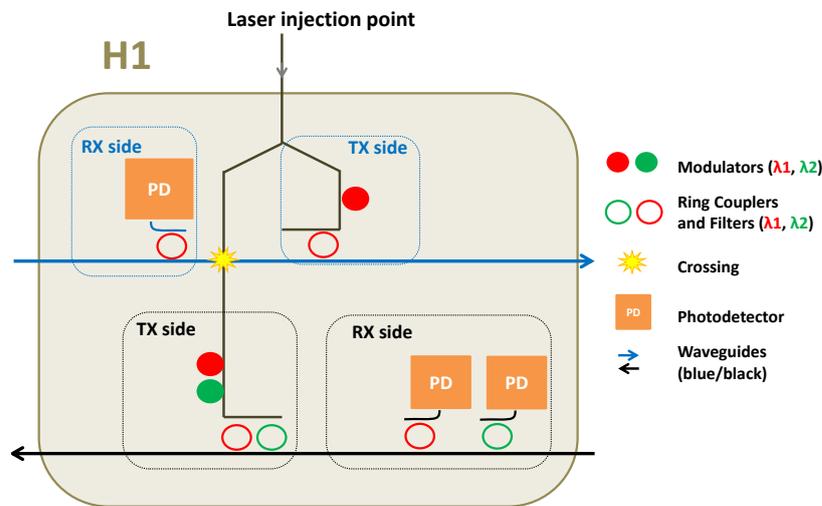


Figure 4.12: Hub architecture of an Optical Ring with physical awareness

the parallel ring waveguides that are furthest away from the injection point (see Figure.4.12). In contrast, such crossings may not appear at target interfaces, since the output signal of photodetectors might directly leave the optical plane by means of TSVs in the best case. However, MRRs are anyway needed to inject wavelengths into and extract them out of the waveguides. Notice that the logic scheme of any Ring topology is characterized by such obstacles, which may degrade the insertion loss, and as a consequence the total power. In this case,  $IL_{max}$  of the Ring logic scheme is about 4.7 dB, while it gets almost doubled when post-layout results are considered. Here, the wiring length contribution becomes dominant through its propagation loss.

The total power consumption of the two topologies is shown in Figure.4.14. Thanks to the lower insertion loss, the 7-way Ring topology results more efficient than the 8x8 Folded Crossbar by about 30%. The latter one is heavily penalized by the larger number of crossings, and the higher wiring length. In this case, the insertion loss gap of 50% is reduced to 30% in terms of total power, due to the relevant contribution of optical receivers in the breakdown. They contribute for 63% in the crossbar topology and for 89% in the Ring one. Ultimately, the optical Ring is clearly an appealing solution for such a small scale system. However, in highly integrated systems the picture is not entirely clear, since the larger connectivity requirements cause the number

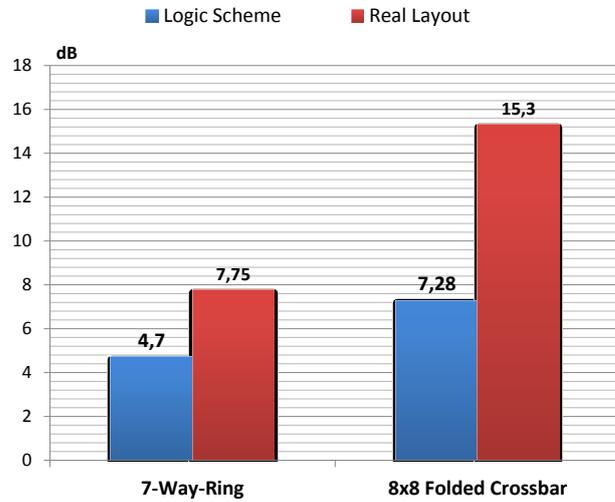


Figure 4.13: ILmax Contrasting: 7-way Ring vs. 8x8 Folded Crossbar

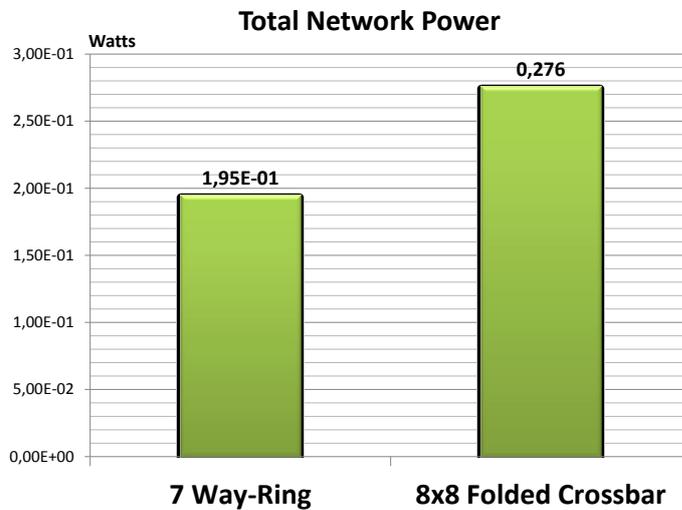


Figure 4.14: Total power Contrasting: 7-way Ring vs. 8x8 Folded Crossbar

of ring waveguides to increase, hence worsening the crossing concerns at initiator interfaces. At the same time, in larger dies the critical path becomes propagation-loss dominated, thus raising another concern for optical Ring networks. Finally, the quality metrics of filter-based topologies will be effectively scaled up in future systems as CAD tools for automatic place&route

of ONoC topologies could become available.

## 6 Network Partitioning

In order to increase the level of confidence of this comparative framework, optimization techniques well beyond global connectivity are worth exploring. In this section in fact we propose Network Partitioning, an alternative design methodology [53] to the global connectivity. Fundamentally, this new design method is advantageous as :

(a) **enables** that wavelengths can be reused across subnetworks similar to what done in telecommunication networks (e.g. laser sources are reused across subnetworks), (b) **simplifies** connectivity patterns, and (c) **allows** that distinct traffic classes can be used across subnetworks. In our design each network partition utilizes 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 (maximum number for global connectivity) to just 4 due to their reuse.

### 6.1 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 (the same ones used for global connectivity, although scaled down), in addition to engineering an ad-hoc topology for the 3D-stacked system at hand.

As mentioned in the previous section.5.1, the **8x8-GWOR** is a scalable and non-blocking wavelength-routed optical router. The basic cell of the former solution is represented by the **4x4-GWOR** that has 4 bidirectional ports located on the cardinal points. Furthermore, two horizontal and two vertical waveguides are used, which intersect each other to form a basic check shape, and MRRs are placed pairwise on waveguide intersections. GWOR does not support self-communication, hence its use for the memory request

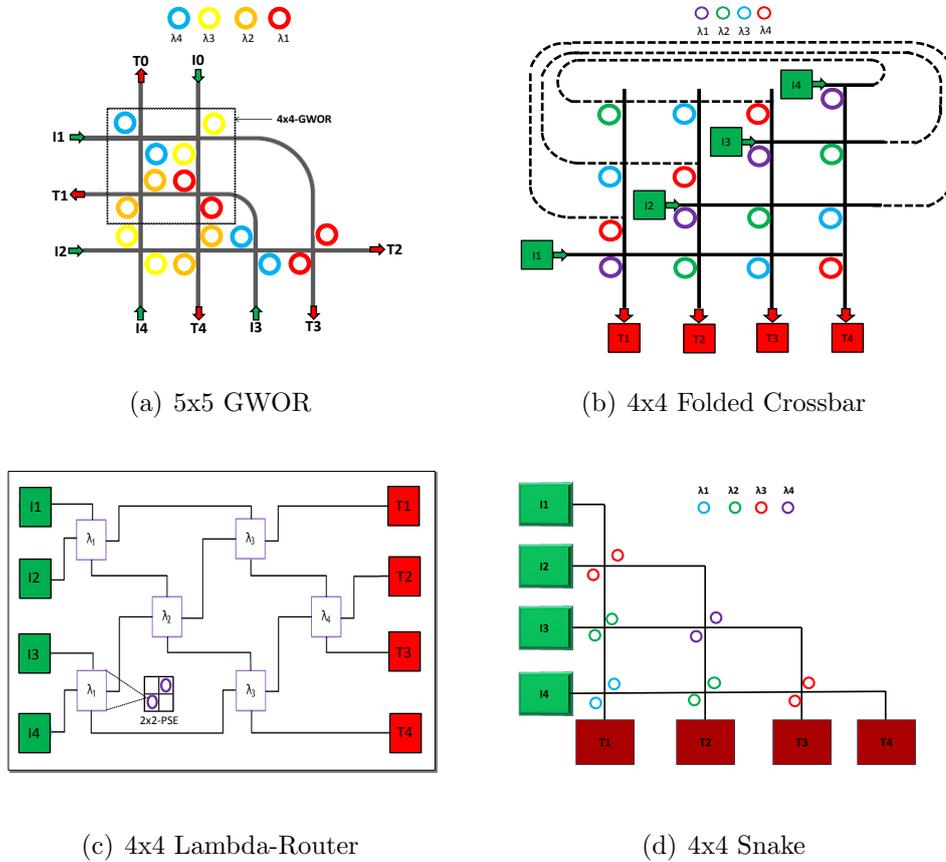


Figure 4.15: Logic schemes of WRONoC topologies under test

and response networks requires its extension to a 5x5 configuration. This is possible, since the wavelength assignment in [28] enables any size of the topology. Figure.4.15(a) shows 5x5-GWOR which it is constructed starting from its lower basic cell (4x4-GWOR). With respect to the baseline scheme, 3 MRRs need to be inserted 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 are removed.

An alternative topology is illustrated in [41] and it is named **4x4 lambda Router**. In order to interconnect 4 masters with 4 slaves, the network makes use of 4 stages of 2 and 1 add-drop optical filters (Figure.4.15(c)). This network is obtained by scaling down the previously proposed **8x8 lambda Router**, or vice versa it could be seen as the preliminary cell to build any size of the lambda router topology. With respect to the original scheme, we replaced the native parallel 2x2 add-drop filters with 2x2 photonic switching

elements, the only difference being an easier physical design thanks to the orthogonally intersected waveguides.

Figure.4.15(b), shows the scaled-down version of the **8x8 Folded Crossbar**, that is customized for connecting 4 initiators with 4 targets (here it is referred to be the **4x4 Folded Crossbar**).

Finally, a custom tailored solution for processor memory communication is described, namely the **Snake** topology. The Snake's pattern (Figure.4.15(d)) is also flexible, since a different number of initiators and targets can be easily accommodated. In the **4x4 Snake**, six 2x2 optical filters are tuned to different wavelengths and their number scales up from the rightmost side to the leftmost one. 4 main optical links connect the slaves while enabling some placement flexibility. This topology is conceived to map efficiently to the placement constraints of the target system.

Ultimately, as it was done for global connectivity, an optical Ring topology is designed to connect 4 masters and 4 slaves, and then it was added in the comparative framework. In practice ring waveguides are devoted to specific message classes in this instance (hereafter it is referred to as **ORNoC**).

For the sake of comparison, all topologies are constrained to use the same number of wavelengths, and laser sources, to instantiate physical resources accordingly. Therefore, all topologies deliver the same overall bandwidth of 440 Gbit/s, and the **4x4 GWOR** is assumed to be the inter-cluster network as its shape better fits the placement constraints of Hubs (positioned along a square in the middle of the optical layer) than other filter-based optical networks.

## 7 Snake vs. **Lambda Router**

With respect to the **Lambda Router**, the proposed Snake topology first breaks the mono-dimensional assumption, and easily fits a planar distribution of initiators and targets. Unlike the **Lambda Router** that grows horizontally, Snake is developed vertically (see Figure.4.17). Also, as shown in Figure.4.16 Snake is easily capable of providing asymmetric solutions (e.g. 4x8, 8x4), much more easily than **Lambda Router**, although this is in principle still feasible with this topology as well. Finally, the Snake topology is engineered

Table 4.3: Layout-aware properties of topologies under test

	Total # of Wavelengths	Max # of Crossings	Max wire length (cm)	Total # of MRRs	Type of MRR
ORNoC	4	3	3.2	40 (8 IC)	4
4x4 SNAKE	4	6	2.4	32 (8 IC)	4
4x4 $\lambda$ -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

to efficiently meet the placement constraints of the target system. In fact it is the custom-tailored solution for the system at hand, and it has the valuable property that connections from the center of the chip to its upper or lower side are much facilitated, thus providing more routing flexibility.

## 8 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. Manually layouts have been drawn with similar criteria to those adopted for global connectivity. The only approximation lies in the lack of the network for the distribution of the optical power. It is as the top-level clock tree the layout of an electronic NoC was neglected.

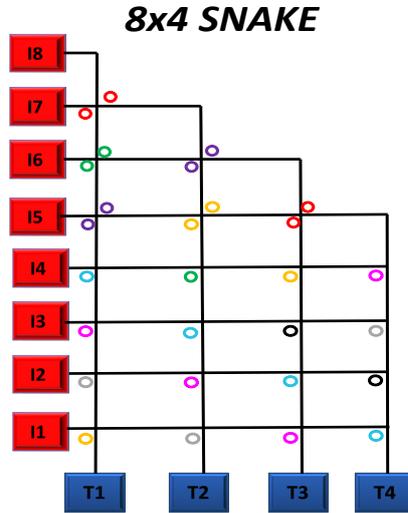


Figure 4.16: Asymmetric 8x4 Snake

After place&route all topologies, the difference between the logic schemes and their layouts is still apparent even if network partitioning mitigates this effect to a significant extent. 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 [55]. **First**, each optical filter is placed close to its attached node; **Second**, filters without any node connection are homogeneously spread throughout the 2D floorplan in order to balance length of waveguides, and above all to avoid waveguide crossings as many as possible. Since these latter play a dominant role in determining the minimum optical power that laser sources should provide to satisfy specific detector sensitivities, as made for global connectivity we consider the elliptical taper, technique already implemented for global connectivity and the aggressive optimization based on MMI tapers[61].

5x5-GWOR (Figure.4.18(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, whereas the total number of MRRs achieves 40 (8 in the inter-cluster network, plus both 16 for the memory request network and the response one). Despite a higher worst case number of crossings in the logic scheme (6), the layout of the 4x4 Folded

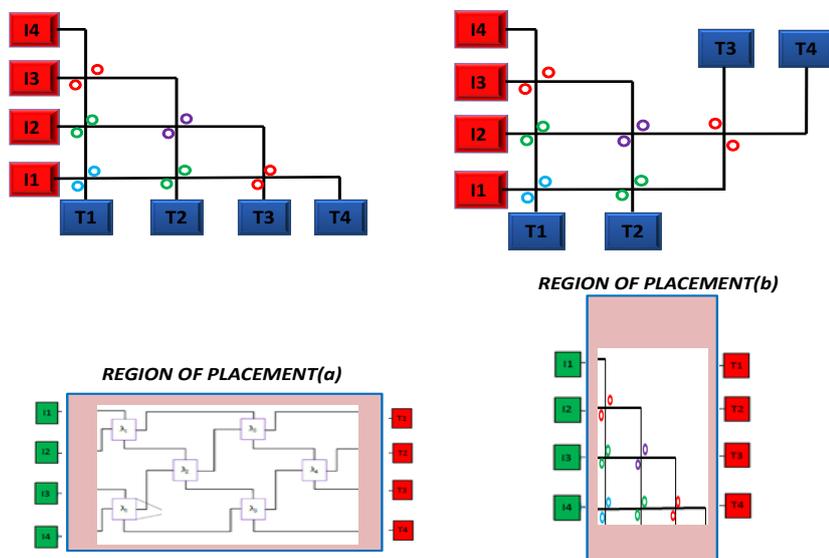


Figure 4.17: Snake Properties

Crossbar Figure.4.18(b) resulted only in 21 crossings, with the same number of MRRs reported in GWOR (40). The layouts of the 4x4 Lambda Router (Figure.4.18(c)), ORNoC (Figure.4.18(d)), and 4x4 Snake (Figure.4.18(e)). are clearly less intricate than the previous ones, hence potentially resulting in lower insertion loss critical paths. More precisely, Lambda Router reports 8 crossings while Snake only 6. By using the wavelength assignment in [56] and a convenient ordering of nodes along waveguides, ORNoC turns out to exhibit 3 crossings on the critical path, all of them are localized inside the network interfaces (HUB) for the sake of waveguide reachability. The Key properties of topologies under test, measured after their physical design, are summarized in table.4.3. 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 has to be used for ORNoC to achieve the same goal. 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 the 8 ones needed in the inter-cluster network. All other topologies instead do not have any injection filters, since

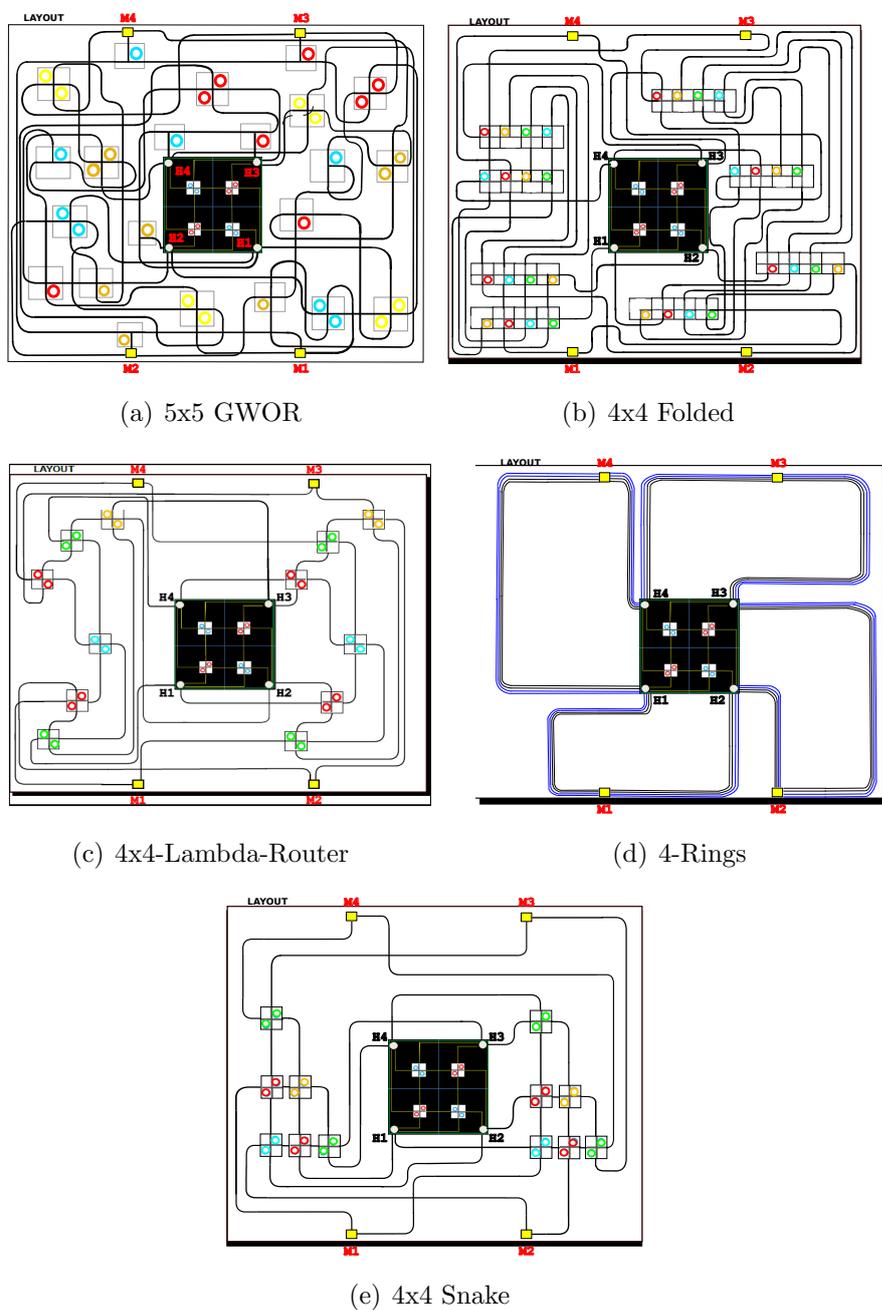


Figure 4.18: Layout of the Optical layer with network partitioning after manual place&route. Requests networks are on the left while response ones on the right of the layout.

they get a branch of the light distribution network which directly enters the network. In the Ring topology, the injection waveguide needs to be bridged to the ring waveguides. Extraction filters at receivers are common for all

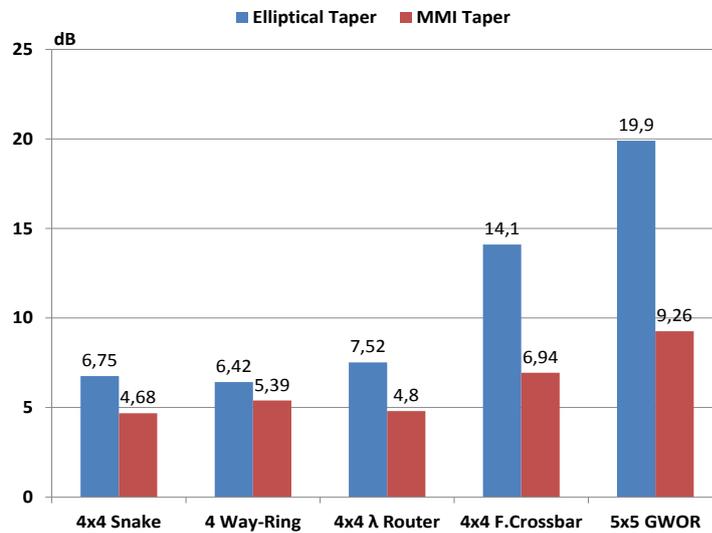


Figure 4.19: Contrasting of the maximum insertion loss across topologies

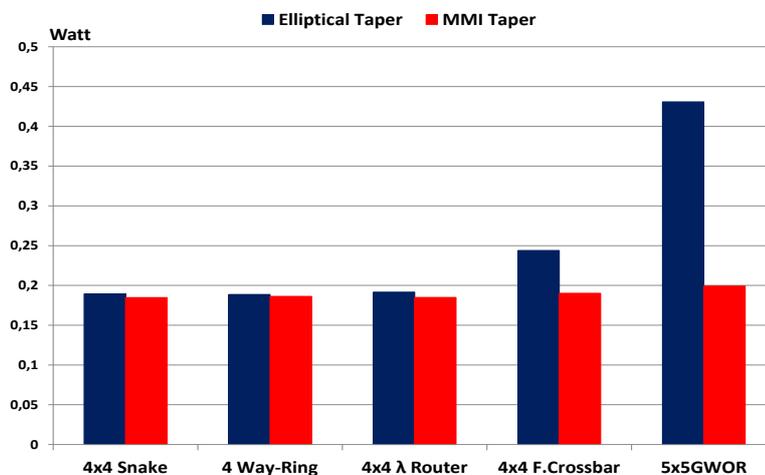


Figure 4.20: Total power comparison across topologies

topologies, hence are not considered in the count.

## 8.1 Power efficiency of topologies

As it was made for global connectivity, our study assumes loss parameters reported in table.4.2. Then, we rely on a Simulink simulation framework to

quantify physical metrics of optical networks. For the evaluation of both insertion loss critical path and total power, we followed the same methodology reported in section.5.2

Figure.4.19 shows the worst-case insertion loss across all topologies considered in this comparison, assuming two kinds of tapers for the optimization of waveguide crossings: the standard elliptical taper (already implemented in the global connectivity scenario) and the efficient Multi-Mode-Interference one (MMI). However, the feasibility of the MMI taper should not be taken for granted, since it depends on the maturity of the manufacturing process, and on the device size. In fact, it has a larger area footprint with respect to the elliptical taper, therefore it might be suitable for layout-induced waveguide crossings, but might be unfeasible for the internal crossings of photonic switching elements, where ring resonators should be placed close to the waveguides for the sake of efficient coupling. Loss parameters utilized for this optimization were derived from 2D-FDTD (Finite-Different-Time-Domain) simulations, and exhibit a crossing loss =0.18 dB while a drop loss = 0.0087 dB. 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 length as always on the critical path. The Snake topology, with its 6 crossings and the same maximum guide length of GWOR, becomes competitive, since propagation losses are not very relevant at this chip size yet. With elliptical taper, the overhead with respect to ORNoC is just 5%. 4x4 Lambda Router has reasonable results in the comparison since it has 22 mm of wiring length and 8 crossings, while the 4x4 Folded Crossbar is better than GWOR for two reasons: lower number of crossings (21), and 4 mm shorter link length on the insertion loss critical path. The effect of MMI is highly beneficial for the Snake, since it minimizes the impact of its crossings over ILmax, 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 apparently higher complexity. In turn, Snake results in a 2.5%, 32.6% and 49.5% lower insertion loss than Lambda-Router, Folded Crossbar and GWOR respectively.

Figure.4.20 shows the total power across all topologies when the energy consumption of the detector is 395fj/bit (or 3.95mW). Power refers to the scenario where the maximum aggregate bandwidth of the network is used (440Gbit/sec with modulation rates of 10Gbit/sec). As can be seen, the total power of GWOR is higher than that of other topologies regardless of the specific taper. With elliptical tapers, GWOR is clearly infeasible under the given place&route constraints. The same holds, for 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 ILmax of the topologies, however the total network power is dominated by receiver power with current technology assumptions (average 75% with Elliptical taper as opposed to 90% with MMI taper), therefore the remaining gap between topologies in Figure.4.19 maps to the total power gap of Figure.4.20 after going through an attenuation factor: just 15mW of difference between Snake (the best) and GWOR (the worst). Of course, different laser sources (e.g., wall-plug laser efficiency) or receiver (e.g., energy) parameters may further widen the gap.

## 9 Global Connectivity vs. Partitioning

This section describes the comparison between the best topologies, both ring-based and filter-based, implemented both for global connectivity and network partitioning. Figure.4.21 shows the total power comparison between the message-class specific optical Ring (from partitioning optimization) and the global Ring. As can be seen, both implementations almost provide the same total power, particularly, the partitioned Ring consumes 188mW, instead the global one 195mW (roughly 3.6% more). This marginal deviation is determined by the laser power. In particular, the global Ring topology features more laser sources than the local Ring (7 vs. 4). Let us denote also that the receiver power results to be the most important contributor in the power breakdown.

In contrast, Figure.4.22 illustrates the total power comparison of the best

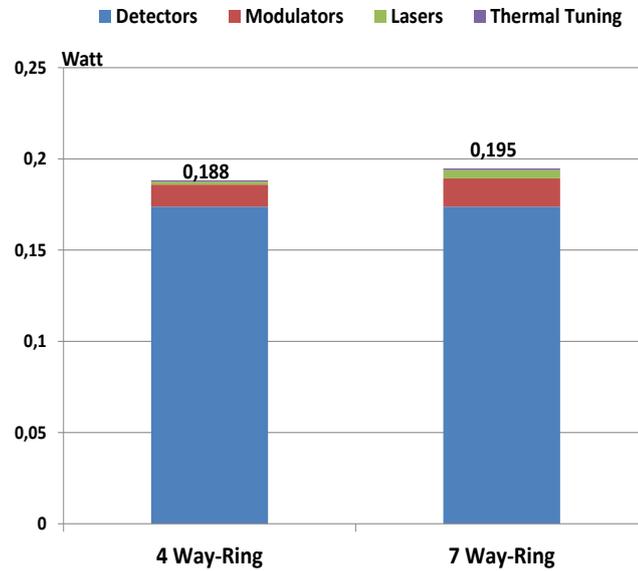


Figure 4.21: Total power comparison: Partitioned vs, Global Ring

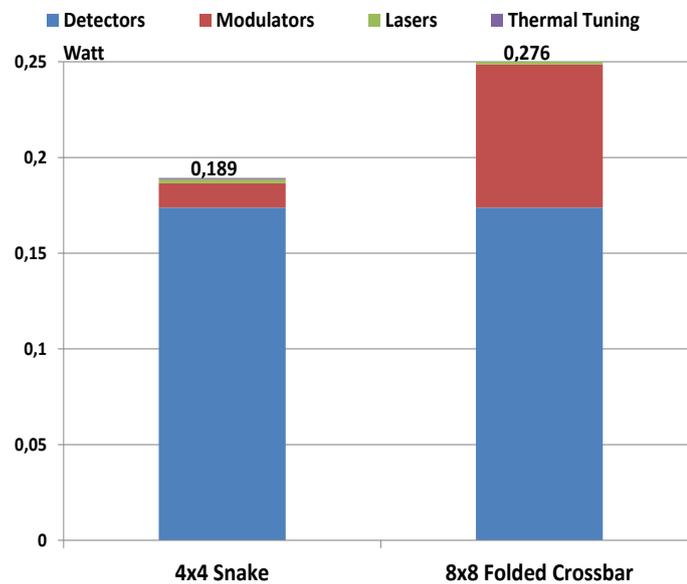


Figure 4.22: Total power comparison: Partitioned vs, Global Ring

filter-based topologies, both in the global connectivity and the network partitioning cases. Similar to ring-based topologies, the receiver power significantly contributes to the total power. However, for the 8x8 Folded Crossbar (the best filter-based global topology) such impact ends up being mitigated by modulator power, that becomes relevant since the 8x8 Folded Crossbar topology provides higher insertion loss than the 4x4 Snake (15.3 dB vs. 6.75 dB). It should be recalled that modulator power partly depends on the input optical power. By including all contributions, the 4x4 Snake results to be more efficient than the 8x8 Folded Crossbar by about 30% (189 mW vs. 276mW). From the proposed analysis, it is clear that filter-based topologies benefit the most from the network partitioning optimization, while ring topologies are in any case global structures. Also, in small-scale systems topology selection can be ultimately dictated by design simplicity considerations, since it is not difficult to engineer both filter-based and ring-based topologies with similar power figures. From this viewpoint, ring-based solutions are clearly appealing.

## 10 Scalability Implications

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 [30]. Consequently, the die sizes grow to 16 mm x 16 mm. We limit the comparison between ORNoC and the best filter-based 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 [27] while conservatively keeping the same sensitivity, a projection which is supported by the physical considerations in [67] about

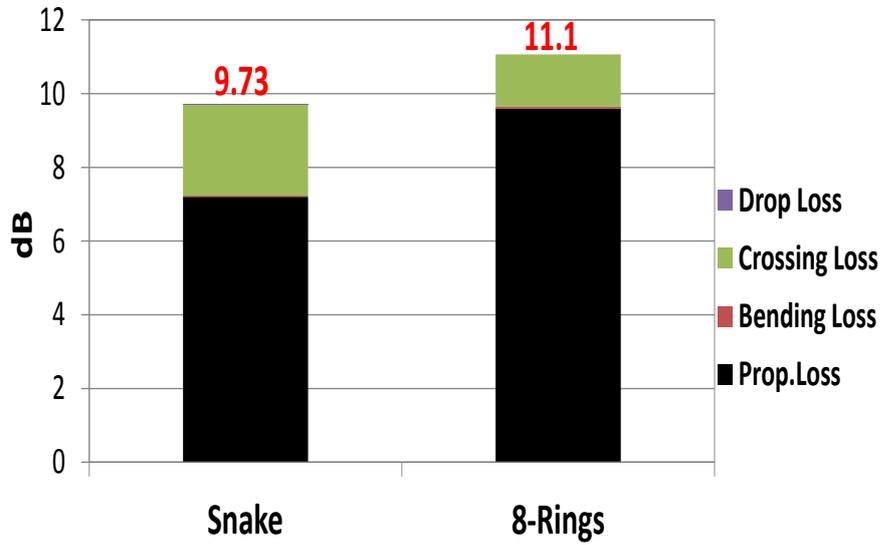


Figure 4.23: ILmax under Scaled Assumptions: Snake vs. Rings

silicon photonics in 3D-stacked systems and receiver circuitry.

Figure.4.23 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.4.24. 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.

## 10.1 System-Level Implications

In section 8.1 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

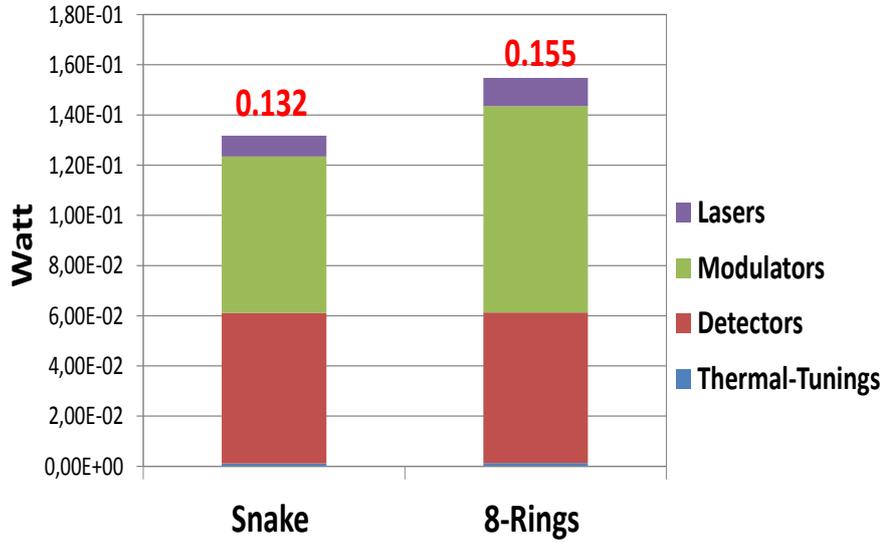


Figure 4.24: Total Power under Scaled Assumptions: Snake vs. Rings

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 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 [76], in which we model the clustered 16-core architecture described in Table 4.4 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 [83]. Performance were evaluated for the Parsec 2.1 multithreaded benchmark suite [77], 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 4.25 shows the performance improvements that can be achieved at system level when different degrees of broadband switching are used and under

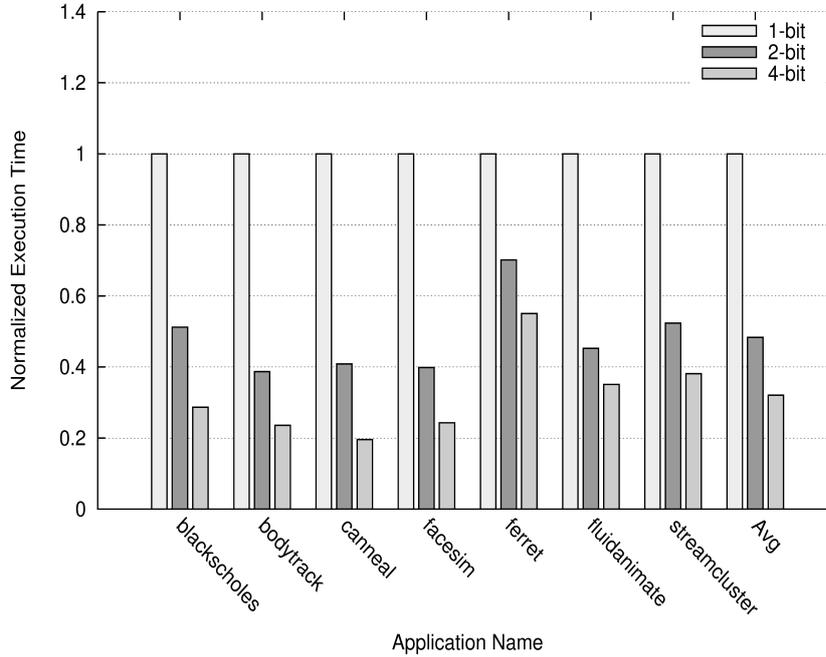


Figure 4.25: System-level performance speedup (normalized).

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 work.

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.

Table 4.4: 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

## 11 Conclusion

This chapter argued 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. **First**, there is a large deviation of insertion loss between the logic scheme and the physical implementation as an effect of placement constraints. **Second**, the most promising logic schemes may turn out to be the worst physical topologies, so the design predictability gap should be carefully quantified. **Third**, network partitioning is an effective way of reusing wavelengths and simplifying ONoC design. **Fourth**, the best topologies for global connectivity are not necessarily the best options for network partitioning. On one hand, for small scale systems, a Spatial-Division-Multiplexed Ring topology is hard to beat. Even in this context, should technology evolutions improve optical receiver energy, filter-based networks could again have a role. In practice, an optical Ring is ideally the best WRONoC topology, but its practical nonidealities (e.g., waveguide reachability in the injection system, worse waveguide length scalability on the critical path) make an actual comparative test with other filter-based topologies mandatory in the target system. On the other hand, for future larger scale systems, where connectivity requirements and die size increase, Spatial-Division-Multiplexing combined with the relevant role of propagation losses seriously penalizes optical ring architectures, so that filter-based topologies may become appealing. This trend will be further consolidated by the devel-

---

opment of CAD tools for automatic place&route of filter-based topologies, which will optimize their quality metrics in layout-intricate and/or highly-integrated scenarios



# Chapter 5

## Network-Interface Architecture for Wavelength-Routed Optical NoC Topologies

### 1 Abstract

This chapter focuses on the description of the network interface architecture for wavelength-routed optical network-on-chip. Figure 5.1 shows the detailed scheme of the proposed network interface. Although the chapter 4 has introduced the key role of the network interfaces for wavelength-routed ONoC, it has not discussed important details and architectural considerations. In contrast, the objective of this chapter is not to present the best possible design point, but rather to start considering the basic components, and indicating which one deserves the most intensive optimization effort for prime time of optical interconnect technology in industry. In addition, the optical NoCs move most from their control logic to the network interfaces (NIs), which should not be oversimplified with abstract models.

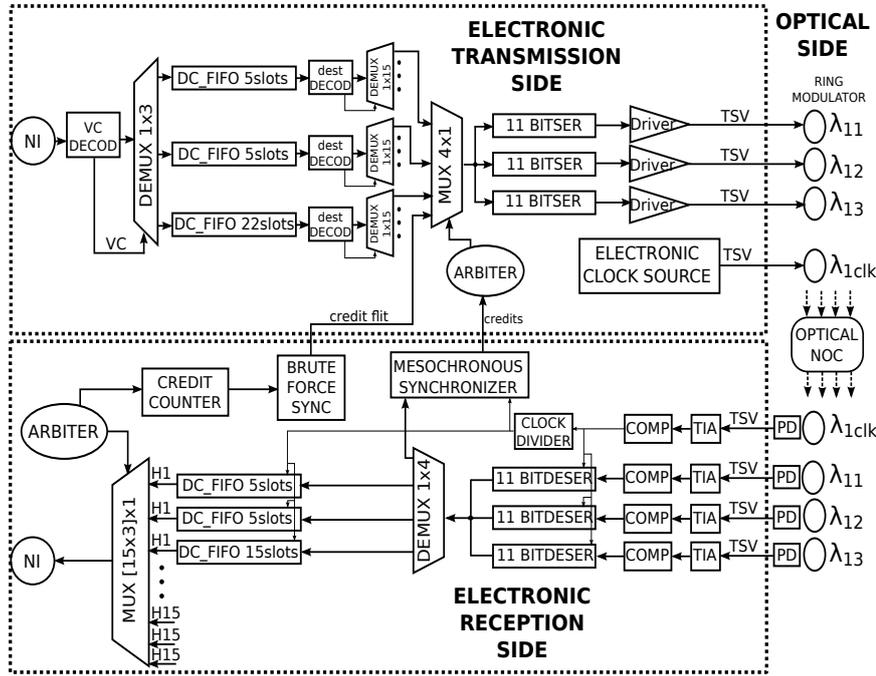


Figure 5.1: Optical Network Interface Architecture.

## 2 Network Interface Architecture: A More Detailed View

Here, we accurately describe the functionalities of most the important parts comprising the proposed Network Interface Architecture.

### 2.1 Wavelength Routed NoC

As widely discussed in chapter.2, the Wavelength-routed optical NoCs rely on the principle of wavelength selective routing. As illustrated in Figure.5.2, conceptually, every initiator can communicate with every target at the same time using different wavelengths. For instance, the first initiator uses wavelengths 1, 2, 3, and 4 to reach 4 different targets 1, 2, 3, and 4, respectively. The topology connectivity pattern is chosen to ensure that wavelengths will never interfere with each other on the network optical paths. This way, all initiators can communicate with the same target by using differentiated wavelengths. WRONoCs support contention-free all-to-all communication with a modulation speed of 10 Gbps/wavelength. The proposed Network Interface

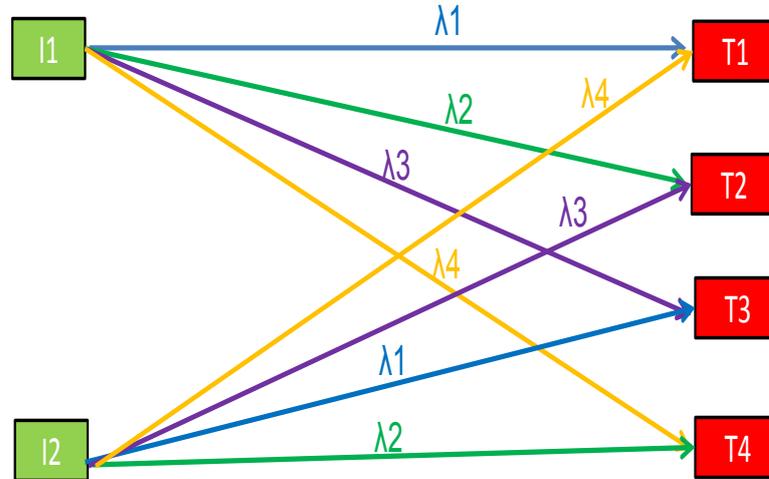


Figure 5.2: Principle of the Wavelength Selective Routing.

Architecture can work with any WRONoC topology.

## 2.2 Message Dependent Deadlock Avoidance

Message Dependent Deadlock (MDD) arises from the interactions and dependencies created at network endpoints between different message types [80, 70]. Figure.5.3 shows the dependence between a request and response at the NI. In a complete system, the combination of these effects may lead to cyclic dependencies. Message dependent deadlocks, once they occur, block resources at both network endpoints and inside the network indefinitely, even if an algorithm is used to avoid routing-dependent deadlocks in the network-on-chip. This arises from the fact that network routers are unable to differentiate between message-dependent deadlocks and normal network congestion. When we apply these considerations to WRONoCs, the problem gets simplified by the fact that there is no buffering inside the network. Therefore, the ONoC automatically satisfies the consumption assumption, which is a necessary (but not sufficient) condition for deadlock avoidance. To enforce the sufficient condition, we must allocate a different buffer for each kind of message in the NI. This has direct implications on the buffering architecture of our target NI (that is, on the number of virtual channels), depending on the communication protocol the WRONoC needs to support. To avoid message-dependent deadlock, every network interface needs separate buffer-

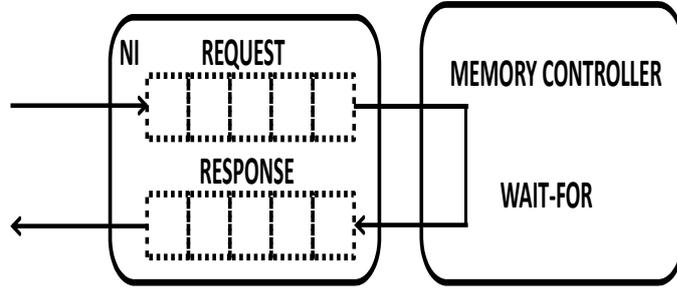


Figure 5.3: Dependence between a Request and Response at the NI.

ing resources for each one of the three message classes of the MESI protocol. This should be combined with the requirements of wavelength routing: each initiator needs an output for each possible target, and each target needs an input for each possible source. As a result, in the baseline version of the NI, each initiator comes with 3 FIFOs for each potential target, and each target with 3 FIFOs for each potential initiator. In a more optimized version of the NI (the one in Figure 5.1), all destinations share the same set of 3 FIFOs and the flits are sent to different paths afterwards (all logic components after 1x15 demuxes are replicated for each destination).

### 2.3 Buffering Sources

All the FIFOs at both the transmission and the reception side must be dual-clock FIFOs (DC FIFOs) to move data between the processor frequency domain ( it is assumed to be 1.2 GHz) and the one used inside the NI. In this work we utilize the DC FIFO architecture proposed in [79]. These devices depend on the bit parallelism. The size of each DC FIFO is chosen based on corresponding size of the packets that use each of the Virtual Channels (VCs). Since control packets need 2 flits, while data packets 21 flits, assuming a flit width of 32 bits, the minimum size to achieve the perfect throughput is 5-slots [79], hence all the VCs at the transmission side are size accordingly. In order to avoid that communications are interrupted, at the reception side the data VC are sized based on the round-trip latency. Hence, at the reception side, the DC FIFOs for data packets are result 15-slots DC FIFOs. On the contrary, for control packets we kept the minimum size (5-slots) since they can already fit two complete packets.

## 2.4 Serialization and Deserialization Procedure

The use of serializers and deserializers is mandatory since it is hardly feasible (e.g. area, power consumption and floorplanning) to integrate as many lasers as the number of bits to transmit per packet (32 bits). For this reason, after flits are forwarded to the appropriate path depending on their destination they need to be converted at 10 GHz in order to be transmitted in the optical NoC. Hence, serializers are used for translating the flit into a 10 GHz bit stream. The number of serializers is defined based on the optical bit parallelism (e.g. 3, 4, etc). In fact 3-bit parallelism means 3 serializers of 11 bits each that work in parallel to serialize 32 bits of a given flit, meaning an overall bandwidth of 30 Gbps. The bit parallelism also determines the frequency inside the optical NI. At 3-bit parallelism, 1.1 ns ( $0.1 \text{ ns} \times 11 \text{ bits}$ ) are necessary to serialize 11 bits while only 0.8 ns are needed with 4 bit-parallelism. The reception side is specular: flits must follow the deserialization process and another set of dual-clock FIFOs.

## 2.5 Resynchronisation: Source Synchronous Communication

Another key issue to be considered in NIs concerns the resynchronization of received optical pulses with the clock signal of the electronic receiver. In this work we assume **source-synchronous communication**, which implies that each point-to-point communication requires a strobe signal to be transmitted along with the data on a separate wavelength, and used to correctly sample received data. Optical transmission of clock signals is an active research field: see for instance [81]. This strobe signal is generated starting from the electrical clock of the transmitter, and removes the need for phase-locked loops (PLLs) or delay-locked loops (DLLs). In particular, the **source-synchronous clock** is utilized to drive the de-serializers, and after the clock divider, the front-end of the DC FIFOs. In this work, we assume that a form of clock gating is implemented, therefore when no data is transmitted, the optical clock signal is gated.

## 2.6 Backpressure Mechanism: The case of the Credit-based Flow Control

Another typically overlooked issue consists of the backpressure mechanism. In this work, we opt for the **credit-based flow control** because it does not rely on timing assumptions, and credit tokens can reuse the existing communication paths, thus avoiding any additional waveguide, and resulting in a milder impact over static power. In addition, the low dynamic energy of ONoC can easily tolerate the overhead of this flow control strategy. All credits are generated at the reception side of the NI when a flit leaves the DC FIFO (at the processor frequency), and forwarded to the transmission side, so that they can be sent back to the source (at the NI frequency). In order to change from one frequency domain to another, we opt for synchronizing the valid bits with a brute force synchronizer. In order to avoid starvation between VCs, when credits arrive to the transmission side, they have the priority over the flits from the VCs. Moreover, when credits arrive at the reception side of the source NI, they need to go through a mesochronous synchronizer to adapt the frequency derived from the received clock to the local NI frequency. Dedicated FIFOs for each source are needed at the reception side of the NIs to support this credit-based flow control. This is a clear candidate for future optimizations.

## 2.7 E/O and O/E Conversions

Once all flits are serialized into bits at 10 GHz such an electrical signal is converted into the corresponding optical one by means of the couple driver-modulator (typically a ring modulator see [30]). After modulation, (**OOK** (On-Off Keying)), the signal can propagate into the optical NoC, and then reaches the appropriate destination based on the value of the physical wavelength. At the reception side, by means of the couple ring filter and photodetector [30], the optical signal is selected and back converted into the electrical one. Finally, it will be collected by a private deserializer after going through the TIA (Trans Impedance Amplifier) and the digital comparator.

### 3 Evaluation

This section shows the most important network-quality metrics for the electro-optical NI: latency, static power, energy-per-bit. Results for an ENoC configured with typical parameters from [68] are also included. This aims to set the bases for a future comprehensive crossbenchmarking study, which is out of the scope of this chapter and it will be pursued in the next chapter6.

Table 5.1: Static Power and Dynamic Energy of Electronic and Optical Devices.

HARDWARE COMPONENTS	3-bit parall			4-bit parall		
	count per NI	POWER (mW)	ENERGY (fJ/bit)	count per NI	POWER (mW)	ENERGY (fJ/bit)
DC_FIFO 5slots (TX)	3	0.12	10.65	3	0.12	12.72
DC_FIFO 5slots (RX)	30	0.12	8.54	30	0.12	10.2
DC_FIFO 15-17 slots	15	0.12	26.50	15	0.12	31.65
DEMUX1x3	1	0.000725	0.92	1	0.000725	0.92
DEMUX1x15	3	0.0021	25.21	3	0.0021	25.21
DEMUX1x4	15	0.00056	6.72	15	0.00056	6.72
MUX4x1 + ARB	15	0.08	0.36	15	0.11	0.49
MUX45x1 + ARB	1	0.9	5.09	1	0.9	5.09
SER	45	0.0475	9.41	60	0.0417	2.63
DESER	45	0.0289	7.74	60	0.0281	6.12
MESO-SYNCH	45	0.041	8.00	45	0.0565	11.1
COUNTER 2bits	45	0.01482	1.014	45	0.01482	1.014
BRUTE FORCE SYNC	15	0.004234	1.4	15	0.00503	1.66
CLOCK DIVIDER	15	0.01172	0.6	15	0.0139	0.714
TSV	120	/	2.50	150	/	2.50
TRANSMITTER aggr	60	0.025	20	75	0.025	20
TRANSMITTER cons	60	0.100	50	75	0.100	50
RECEIVER aggr	60	0.050	10	75	0.050	10
RECEIVER cons	60	0.150	25	75	0.150	25
THERM-TUN /ring	180	0.020	/	225	0.020	/
LASER POWER aggr	/	0.0421	/	/	0.0525	/
LASER POWER real	/	0.308	/	/	0.385	/
E-SWITCH (3VCs)	/	17.9	193	/	17.9	193

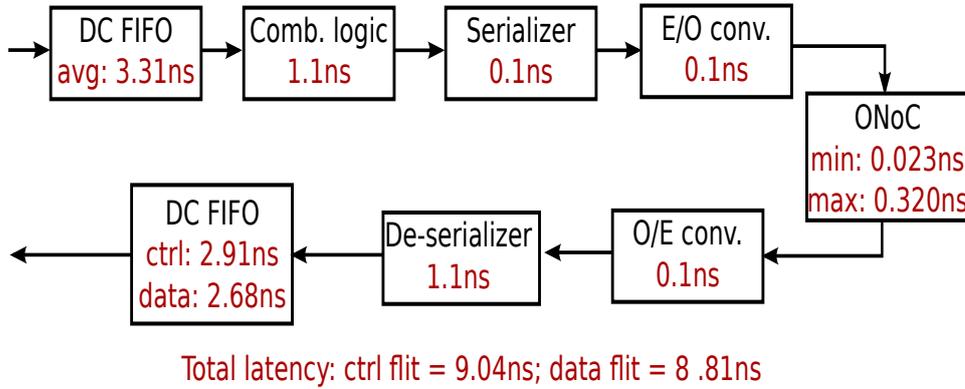


Figure 5.4: Latency breakdown of the optical NI with 3-bit parallelism and the optical Ring.

### 3.1 Methodology

To obtain accurate latency results, we implemented detailed RTL models of the optical and electronic network interfaces and NoCs using SystemC. We instantiated a 4x4 2D mesh for the electronic NoC, and a similar system connected through the optical Ring. The network-wide focus, well beyond the NI, aims at relating NI quality metrics with network ones. Delay values for the optical Ring have been backannotated from physical-layer analysis in [27], and have been differentiated on a per-path basis. For power modeling, every electronic component has been synthesized, placed and routed using a low power 40 nm industrial technology library. Power metrics have been calculated by backannotating the switching activity of block internal nets, and then importing waveforms in the Prime-Time tool. We have applied clock gating to achieve realistic static power values. Energy-per-bit has been computed by assuming 50% switching activity.

Table.5.1 lists the static power and energy-per-bit for all the electronic and optical devices. For the fast developing optical technology, we consider a coherent set of both conservative and aggressive values (obtained from [30]). These values are only realistic under the assumption of low network contention, which rejects the typical operating condition of cache-coherent multicore processors.

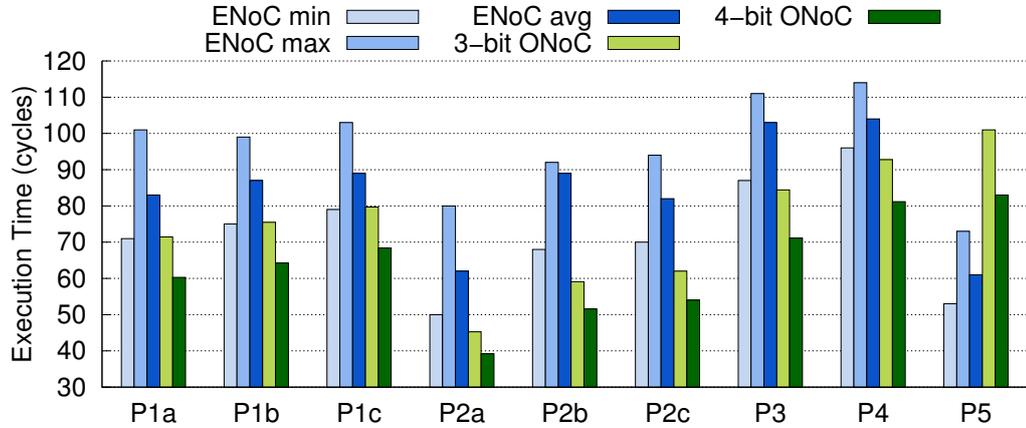


Figure 5.5: Latency of the most common communication patterns. For the ENoC, we include minimum, maximum, and average paths.

### 3.2 Latency Breakdown

Figure.5.4 presents the latency breakdown for the NI components and the ONoC with 3-bit parallelism, obtained from our accurate RTL-equivalent simulations. As can be seen, the latency of the network is negligible (it reports 23 ps and 320 ps across best and worst case paths), but it requires support from a time consuming NI. In fact, inside the NI, the DC FIFOs are the components with the largest latency (see DC FIFOs at the transmission and reception sides).

### 3.3 Transaction Latency

We simulate the most common traffic patterns generated by a MESI coherence protocol in our RTL models without any contention. The increased accuracy of our analysis stems from the fact that our packet injectors and ejectors model actual transactions of the protocol, as well as their interdependencies. Table.5.2 describes the analyzed compound transactions and Figure 5.5 presents the zero-load latency results.

The messages included in these patterns amount to an average 99.9% of the total network traffic, as we observed from full-system simulations of realistic parallel benchmarks from PARSEC and SPLASH2 and multiprogrammed workloads built with SPEC applications (we only exclude communication

Table 5.2: Messages generated by the coherence protocol.

id	Event	Sequence of messages
P1a	L1 miss	<ol style="list-style-type: none"> <li>1. Request from L1 to L2</li> <li>2. Data reply from L2 to L1</li> <li>3. ACK from L1 to L2</li> </ol>
P1b/c	L1 write miss, 1/2 sharers	<ol style="list-style-type: none"> <li>1. Request from L1 to L2</li> <li>2. L2 sends data reply and invalidates 1/2 sharers</li> <li>3. Sharers sends ACK to L1 req.</li> <li>4. ACK from L1 to L2</li> </ol>
P2a	L1 needs upgrade to write	<ol style="list-style-type: none"> <li>1. Request from L1 to L2</li> <li>2. ACK reply from L2 to L1</li> <li>3. ACK from L1 to L2</li> </ol>
P2b/c	L1 needs upgrade to write, 1/2 sharers	<ol style="list-style-type: none"> <li>1. Request from L1 to L2</li> <li>2. ACK reply from L2 to L1 and invalidates 1/2 sharers</li> <li>3. Sharers send ACK to L1 req.</li> <li>4. ACK from L1 to L2</li> </ol>
P3	L1 write miss, another owner	<ol style="list-style-type: none"> <li>1. Request from L1 to L2</li> <li>2. L2 forwards request to owner</li> <li>3. Owner sends data to L1</li> <li>4. ACK from L1 to L2</li> </ol>
P4	L1 read miss, another owner	<ol style="list-style-type: none"> <li>1. Request from L1 to L2</li> <li>2. L2 forwards request to owner</li> <li>3. Owner sends data to L1 and L2</li> <li>4. ACK from L1 to L2</li> </ol>
P5	L1 replacement	<ol style="list-style-type: none"> <li>1. Writeback from L1 to L2</li> <li>2. ACK from L2 to L1</li> </ol>

with the memory controllers). Therefore, they are a very good indicator of the network latency improvements we can expect from the optical network, including its (non-negligible) network interface overhead. We observe that in all the patterns except the last one, the ONoCs either beat or obtain equal results to the ENoC with all path lengths. As opposed to the ENoC, most of the latency of the ONoC is spent in the NI, which is needed to support the low latency optical communication. The tendency changes in pattern 5 because the replacement packet is using a VC designed for control to transmit data, and the smaller FIFO cannot store enough fits to support the round-trip latency. However, this messages are only 7.4% of the total network traffic.

### 3.4 Static Power & Energy-per-Bit

Figure 5.6 depicts the static power and (dynamic) energy-per-bit for the ENoC vs. the 3 and 4-bit parallelism ONoCs. We do not consider ONoCs with less than 3-bit parallelism because the bandwidth of the optical paths would be too low, or ONoCs with more than 4-bit parallelism, because the static power becomes unacceptable (we can see a clear trend in Figure 5.6). We present a breakdown of the contributions of the NIs and NoCs. For the NI, we also separate the electronic components from the optical (and analogic) ones. The optical NoC is solely composed of laser power, so it has no impact on dynamic energy. In computing total power figures, we consider two sets of parameters for optical interconnect technology, corresponding to its high maturity (named aggressive parameters) and to its low maturity (conservative parameters). We observe that the electronic switches dominate the static power, accounting for 95.8% of the total. However, this trend is reversed in the ONoC, with a contribution of only 10.6% and 11.8% for the aggressive technology with 3 and 4-bit parallelism, respectively. It is worth highlighting that most of the static power of the electronic components in the NI comes from the DC FIFOs.

For energy-per-bit we included minimum, maximum and average path lengths for the ENoC and specific values for control and data packets for the ONoC (which change due to the different size of the reception DC FIFOs). As can be seen the ONoC has significantly lower energy-per-bit than the ENoC. Apart from that, we still see how the main contributor for the ENoC energy is the

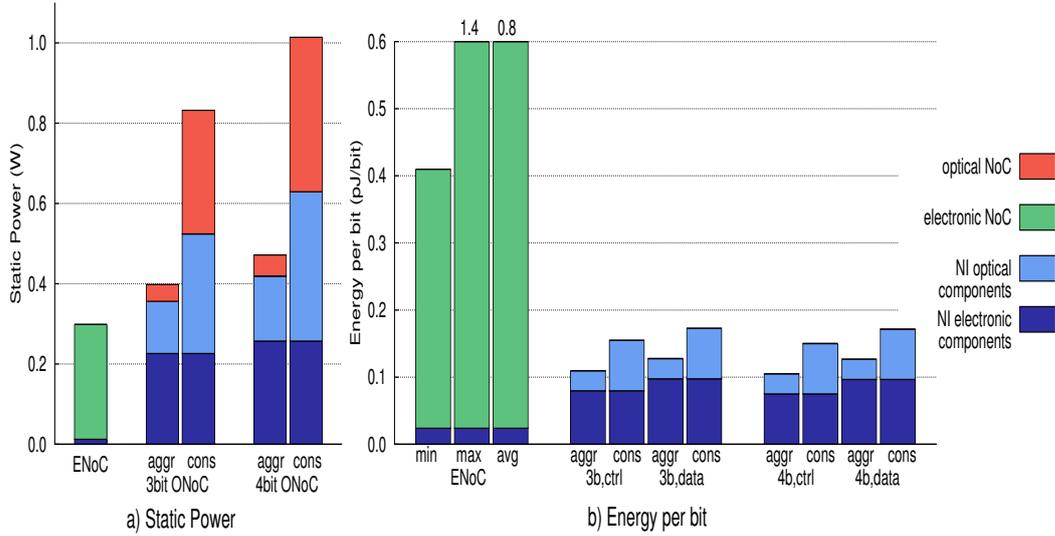


Figure 5.6: Static power and Energy-per-Bit of the NIs and the electronic and optical NoCs.

NoC, while for the ONoC, the complexity is all concentrated inside the NI.

## 4 Conclusion

This chapter presented an accurate design of NIs for WRONoCs and has quantified the effect on the most important network-quality metrics and sets the scene for further improvements of comparative ONoC analysis. Regarding latency, the ONoC is always faster than its electronic counterpart even considering the NI, thus preserving the primary goal of a WRONoC. The behaviour under contention depends mainly on the available bandwidth of the interconnect technologies under test. For the WRONoC, such bandwidth can be modulated by tuning the bit parallelism.

When we consider power figures, we note that while switches are the main contributors in ENoCs, the NI has the largest share in ONoCs. For static power, this contribution is in the same order of magnitude than that from laser sources, for conservative optical technology parameters. However, by further improving the optical technology, the role of the NI becomes dominant, thus making it the main target for future optimizations. Finally, the

---

ONoC preserves its superior dynamic power properties over its ENoC counterpart, even in the presence of its NI.

This chapter has shown that the NI architecture should not be overlooked for realistic ONoC assessments, and comes up with new insights not provided by earlier photonic network evaluations. The most important one is that NI optimizations perhaps have higher priority over the relentless search for ultra-low-loss optical devices. This chapter has also included contents that are referred to a cooperative and interdisciplinary work where further details are in [82].



# Chapter 6

## Crossbenchmarking Framework between the Most Efficient ONoC and its Aggressive Electrical Baseline

### 1 Abstract

Many crossbenchmarking results reported in the open literature raise optimistic expectations on the use of optical networks-on-chip (ONoCs) for high-performance and low-power on-chip communication. However, most of those previous works ultimately fail to make a compelling case for chip-level nanophotonic NoCs, especially for the lack of aggressive electronic baselines (ENoC), and the poor accuracy in physical- and architecture-layer analysis of the ONoC. This chapter aims at providing a crossbenchmarking framework between an optical Network-on-Chip and its aggressive electrical counterpart with realistic complexity, performance, and power figures, synthesized on an industrial 40nm low-power technology. At the same time, key physical design issues for the ONoC under test are carefully assessed, thus paving the way for a well-grounded definition of the requirements for the emerging ONoC technology to potentially achieve energy break-even point with respect to pure electronic interconnect solutions in future multi- and many-core systems.

## 2 Introduction

This chapter aims at a higher level of practical relevance in assessing the potentials of ONoCs for future multi- and many-core systems. This is fundamentally achieved in two ways. On one hand, *we make use of an aggressive electrical baseline*. In fact, realistic design point for the ENoC architecture in terms of complexity, and power are considered. Moreover, real synthesis runs of the target ENoC on a 40nm industrial low-power technology will provide the reference quality metrics the competing optical NoC solutions are contrasted with. On the other hand, *the ONoC is designed and accurately characterized based on both accurate physical-layer and architecture-layer analysis*. The wavelength-routed Ring topology for the ONoC is selected, whose simplicity can reduce the adoption risk of an emerging technology. At the physical layer, the increased accuracy in ONoC modeling is achieved by drawing the Ring layout, especially its injection and ejection interfaces.

At the architecture layer, as previously discussed in chapter.5, the design of the network interface architectures needed to inject/eject electronic packets into/from the ONoC is made, thus capturing typically overlooked sources of performance and power overhead, such as flow control, clock resynchronization, or suitable FIFO sizing.

This work also *carefully considers fixed-power overheads*, which are a significant percentage of total ONoC power. Static power is especially important in those application domains where the network does not undergo high utilization, but it has to serve sporadic traffic peaks. This is the case of shared memory multiprocessors with distributed last-level cache, implementing hardware support for cache coherence. This scenario is becoming mainstream in many application domains, but it is challenging both for the standalone ENoC and for the ONoC. In fact, cache coherence protocols rely on a number of different message types with chained dependencies, thus raising the message-dependent deadlock concern, which is typically addressed by ENoC designers by adding virtual channels. In any case, the use of an ONoC makes sense in this domain only if it can significantly cut down on the total application execution time, thus burning less static power. This chapter considers the case study of a directory-based implementation of the MOESI protocol, and

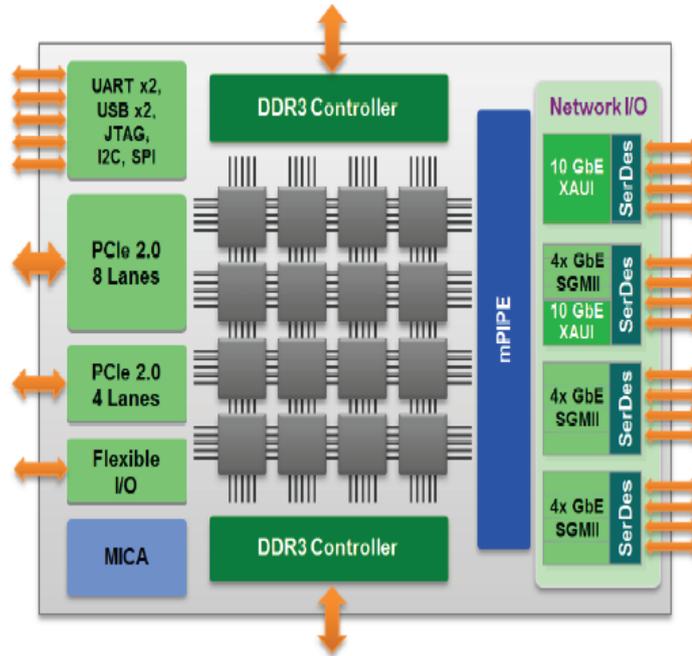


Figure 6.1: Tile 16 (from Tiler Corporation).

derives the requirements for both ENoC and ONoC design.

Finally, the crossbenchmarking effort is extended to the system level by back-annotating the relevant physical and architectural metrics/effects into an abstract system-level functional simulation framework, capable of projecting the ultimate impact that optical interconnect technology may have on realistic execution scenarios. Moreover, realistic traffic patterns such as Parsec 2.1 are used for the estimation of performance and energy metrics for both Electronic and Optical NoCs.

### 3 Target System

Similar to system Tile 16 (see Figure.6.1), our experimental setting consists of a multi-core processor composed of 16 Tiles. Each of them operates as both initiator and target for communications over the system interconnect.

Tiles are disposed over the die area in a common 4x4 2D-Mesh structure (see gray tiles in Figure.6.1). Each Tile is associated with its respective Network Interface (a master or a slave one), performing protocol conversion and (de-)packetization. In our study we assume that these may be omitted during power analysis since they are equally required for both NoC implementations, and will not be cause of differentiation. Our analysis begins from the point where the two architectures start to diverge, including for instance the buffering architecture and frequency converters.

## 4 Baseline Electronic NoC

In this section we introduce the baseline Electronic Network-on-Chip (ENoC). We implement as an aggressive low-power 2D mesh for our reference technology and chip multiprocessor architecture. We chose a 2D mesh topology because it is a regular structure that maps well to the regularity of chip multiprocessor architectures and it is well suited for general-purpose multi-core systems.

The switch architecture is inspired by the  $\times$ pipesLite architecture [68], which represents an ultra-low complexity design point in the space of electronic NoCs.  $\times$ pipesLite architecture is an input-buffered switch, implementing logic-based distributed routing and wormhole switching. In this design, **one clock cycle** is taken to traverse the switch and **one clock cycle** to traverse the link connecting two switches. Buffer capacity is set to two slots for input buffers and six slots for output buffers. The flit width is **32 bits**, which represents a good trade-off between area occupancy and provided bandwidth. Also, a directory-based implementation of the MOESI protocol is considered, which requires at least 3 virtual channels to avoid message-dependent deadlock [70]. Virtual Channel flow control is implemented by replicating the basic switch architecture without VC capabilities three times, based on the approach presented in [71]. This technique results in higher maximum operating speed, better performance and smaller area. In order to preserve the generality of the design and support cores with different operating frequencies that access a fixed-frequency NoC, dual-clock FIFOs are included at the network interfaces.

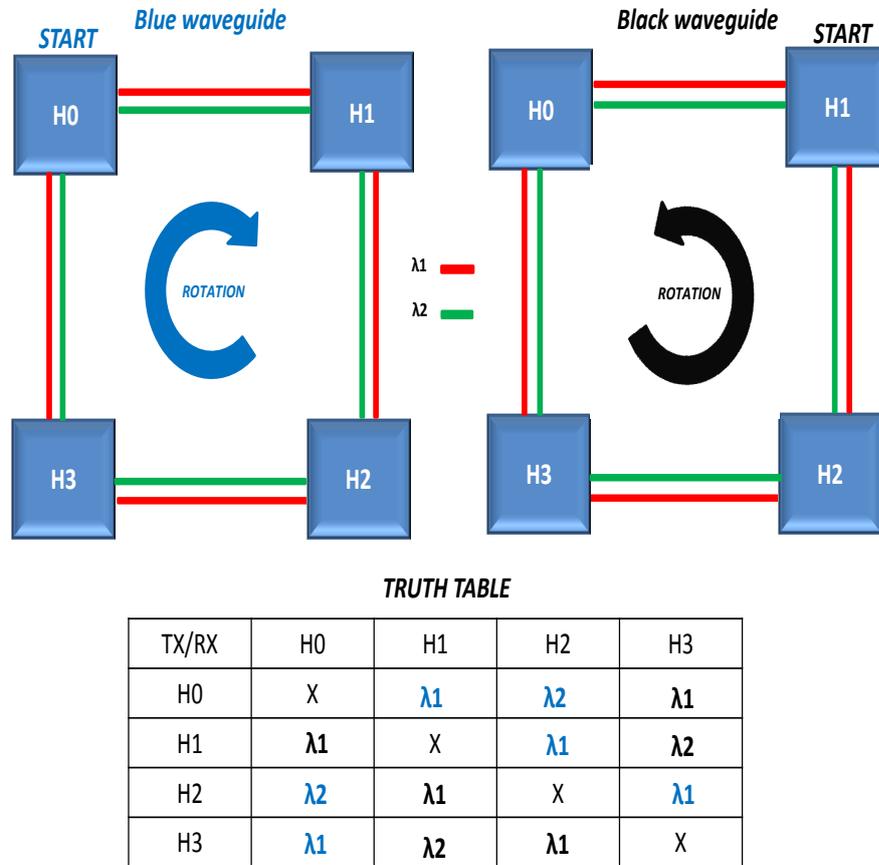


Figure 6.2: Principle of the designed Optical Ring Architecture.

Post-layout analysis are performed to assess performance and power metrics of the Electronic NoC. For this, the switch architecture is synthesized, placed and routed using *Synopsys Design Compiler* and *IC Compiler* tools. For the physical implementation we leveraged on an Ultra-Low-Power Standard VTh 40nm industrial technology library. After layout generation, the maximum operative frequency of our design is considered to be 1.2 GHz.

## 5 Wavelength-Routed Optical Ring Design

In this section we describe our Wavelength-Routed optical Ring topology, relying on a rigorous design methodology, addressing waveguide crossing concerns, and assessing laser power.

## 5.1 Design Methodology

Simplicity and low implementation cost make the optical Ring topology one of the most appealing interconnection networks proposed in the open literature to interconnect initiators and targets of a given multi-processor system-on-chip (MPSoC). In a 3D-stacked scenario, and especially when the global connectivity is required, the optical Ring is certainly the topology that efficiently meets the place & route constraints unlike other solutions such as multi-stage networks, and filter-based ones as widely discussed in chapter 4. For these reasons, we engineer a Wavelength-Routed Optical Ring Architecture by following the principle illustrated in Figure 6.2. The key property is that the same wavelengths can be reused on a single waveguide to establish multiple communications.

Here, we have an optical Ring network structured into 4 **Hubs** ( $H0, H1, H2, H3$ ) which are both initiators and targets. In the proposed example, two different wavelengths ( $\lambda_1, \lambda_2$ ) and two distinct waveguides (the blue and black one) are sufficient to realize 12 contention-free optical paths. This is demonstrated through the wavelengths assignment reported in the truth-table of Figure 6.2. We denote that, no wavelength is listed along its diagonal as self-communications are not allowed. Each wavelength may be also propagated both on the blue and the black waveguides based on two conditions: starting node and rotation. For the blue waveguide  **$H0$**  is assumed to be the starting node while rotation is considered clockwise. On the contrary, for the black waveguide,  **$H1$**  is assumed the starting node while rotation is counterclockwise (see leftmost and rightmost side of the figure 6.2). Once all of two wavelengths are utilized in the first waveguide to populate all of possible contention free optical paths started from  $H0$ , another physical waveguide should be added to populate remaining paths. For this purpose the black waveguide is added. For example, based on above wavelength assignment, the optical signal that resonates at  $\lambda_1$  is used to establish all short range communications among actors (e.g  $H0-H1, H1-H2,..$  etc) which correspond to **one-hop optical path**. While,  $\lambda_2$  ( $H0-H2, H2-H0$ ) are used to accomplish long range distance, with their **two-hops optical path** respectively. By pursuing the presented principle we scaled our Ring architecture up to **16 Hubs**.

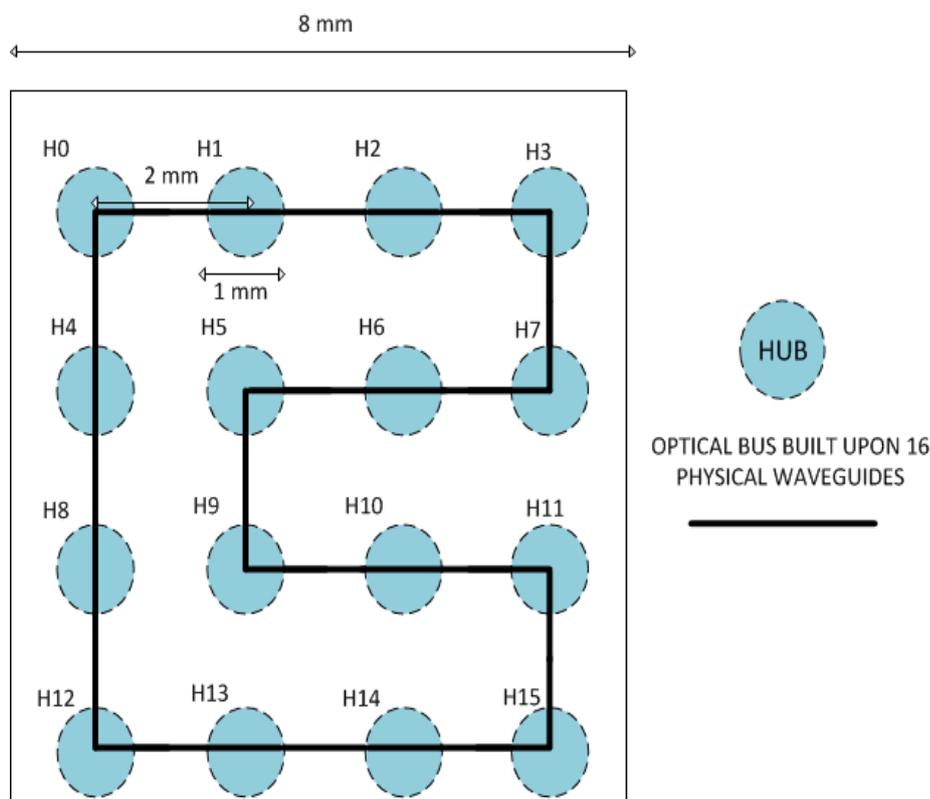


Figure 6.3: Floorplan of a 16x16 Wavelength-Routed Optical Ring Architecture.

In the design of a 16x16 optical Ring, 13 wavelengths are reused and multiplexed on 16 different waveguides to enable 240 contention free optical paths. Figure 6.3 shows the proposed floorplan of our Ring implementation which basically works like a bus since includes multiple waveguides into it. The die size is assumed to be  $8\text{mm} \times 8\text{mm}$ , while inter-hub distance and hub width are assumed  $2$  and  $1$  mm respectively. Ultimately, as it was described here, by leveraging on both SDM (Spatial-Division-Multiplexing) technique, based on replicating physical waveguides, and the WDM (Wavelength-Division-Multiplexing) one which enables several optical communications on a single waveguide at the same time, and alternating also different rotations (clockwise and counterclockwise), it was possible to deliver multiple contention-free optical paths, while minimizing the amount of physical resources and of optical power losses.

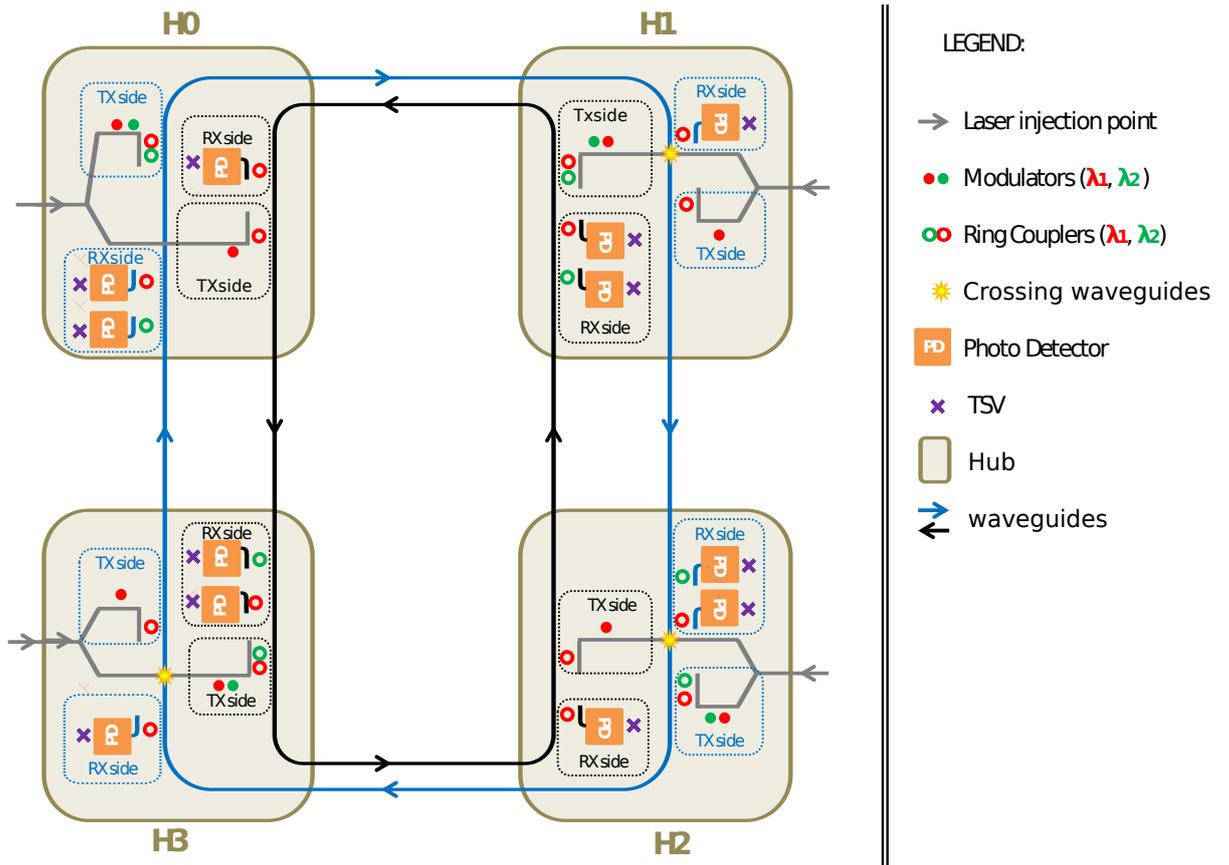


Figure 6.4: 2x2 Optical Ring Topology with physical awareness.

## 5.2 The Waveguide Crossings Concern in Optical Ring Design

In principle in any optical Ring topology neither waveguide crossings nor optical switching elements appear. For this reason, it certainly results an appealing interconnection network. However, there are physical effects that come into play when its actual implementation is pursued. As depicted in Figure. 6.4, the light emitted by laser sources (assumed off-chip) is first multiplexed into a private waveguide (gray arrow line that enters each Hub) and then spread to all optical modulators that are located inside each Hub. As a consequence, the reachability of all waveguides from the injection to the ejection interfaces of optical packets has to pay the price of undesired crossings (see yellow star). Furthermore, MRRs (Micro-Ring-Resonators) are needed not only at the destination stage (Ring Filters) to selectively eject the optical signal, but also to couple the optical same signal into the Ring network after modulation (Ring Couplers).

It is reasonable to expect that such physical effects become even more important when the Ring topology is scaled up to 16 Hubs. In particular, the worst case number of crossings inside the optical network interface is 15 and the total number of MRRs is 2.16K, which means more insertion loss and thermal tuning overhead .

## 5.3 Laser Power Assessment

The preliminary step to evaluate the efficiency of an optical network relies on the estimation of the insertion loss across all wavelengths involved in the given design. Such a metric is extremely important to quantify the total amount of laser power needed to reliably detect the optical message at the destination node. For this reason, we calculate the insertion loss (IL) as the sum of physical components which affect the optical signal along the path, starting from modulator, coupling filter, propagation distance, bending waveguide, ring filter, photodetector without overlooking crossing waveguides. In the end, we quantify the worst case  $IL_{max}$  on each wavelength and laser power is derived accordingly. In our study, die size is assumed to be  $8mm \times 8mm$  while loss parameters are listed in table 6.1.

For the sake of a more comprehensive analysis of nanophotonic devices, and of their evolution over time, we distinguish two relevant cases: realistic and aggressive ones. For the former one we consider a wall-plug laser efficiency of 8% while crossing waveguides are supposed to be optimized with the standard elliptical taper [60]. In the aggressive case, these quality metrics are projected assuming a laser efficiency of 20% whereas crossings are improved through MMI tapers [61]. In both cases, the detector sensitivity is considered the same  $S = -20$  dBm.

Figure 6.5 illustrates the laser power trend across wavelengths. This figure highlights that laser sources (assumed to be Continuous Wave) must be treated in a different way. In other words, as the insertion loss is not the same for each path, CW lasers can be sized accordingly. Therefore, there will be some laser sources that turn out to be more power hungry than others. When the aggressive case comes into play, even if the absolute power gap is strongly reduced (4x), there is still a relative gap across wavelengths, thus confirming that laser sources should be treated separately. Finally, had we ideally designed our ring topology without accounting for crossings at each optical interface, the total laser power would have been clearly lower than our accurate approach. More in details, there would have been a reduction of 80%, and 41% in the realistic and aggressive scenarios respectively.

## 6 Power Modeling

In this section we describe the power modeling assumptions on which our crossbenchmarking framework is based. Every electronic component is synthesized, placed and routed using a Low-Power 40nm industrial technology library, in order to provide realistic power measurements (not derived from optimistic or ideal estimations). Power metrics are calculated by backannotating the switching activity of block internal nets, and then importing waveforms in the *PrimeTime* Tool. It is worth observing that clock gating technique is applied for the sake of realistic measurement of static power.

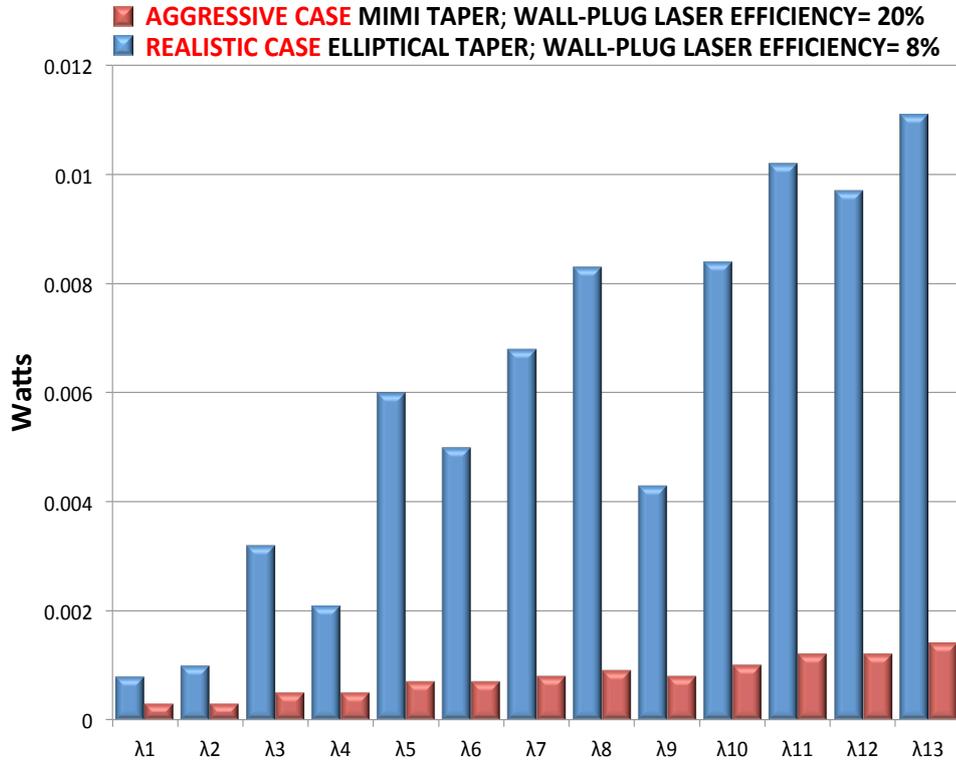


Figure 6.5: Laser power results across wavelengths: aggressive vs. realistic.

PHOTONIC COMPONENTS AND DEVICE PARAMETERS	VALUE
COUPLER LOSS	0.46dB
MODULATOR INSERTION LOSS	4.0dB
PHOTODETECTOR LOSS	1.0dB
FILTER DROP LOSS	1.0dB
THROUGH RING LOSS	$10^{-3}$ dB/ring
PROPAGATION LOSS	1.5dB/cm
BENDING LOSS	0.005dB
WAVEGUIDE CROSSING LOSS (@REALISTIC)	0.52dB
WAVEGUIDE CROSSING LOSS (@AGGRESSIVE)	0.18dB
RECEIVER SENSITIVITY	-20dBm
LASER EFFICIENCY (@REALISTIC)	8%
LASER EFFICIENCY (@AGGRESSIVE)	20%

Table 6.1: Photonic Components and Device Parameters.

Energy-per-Bit is computed by removing the Static Power by the Total power on a component-basis, under 50% switching activity assumption.

Overall, the power consumption of the Electronic NoC is built upon replicating the power contribution of its basic switch components.

As mentioned in section 4, we also considered the power consumption of both electronic NI buffering and frequency converters (dual-clock FIFOs) which contribute around 11.5 mW. The static power dissipated (Idle power) by the **entire electrical network** (16 switches), is around 286 mW (only the top-level clock tree is omitted). In addition, the energy required for transmitting data over each hop of the ENoC is 193 fJ/bit.

Similarly, the power dissipation of **Optical Network Interfaces** is computed by composing the power consumption of each of its sub-blocks such as DC\_FIFOs at the transmission sides, Demultiplexers, SERs, Synchronizers, DESERs, DC\_FIFOs at the reception sides, Multiplexers, and Credit counters, as widely explained in the previous chapter 5.

The static power contribution of all **optical components** is given by: Laser sources, Thermal tuning, Transmitter (i.e., the driver-ring modulator couple), Receiver (i.e., Photodetector, Trans-Impedance Amplifier, and Comparator) and the source-synchronous clock. The latter addendum is internally composed by further laser sources, Transmitters, Receivers, and MRRs as well. For static and dynamic power parameters, as well as for their relative ratio, we consistently assume values from the same literature source [30, 72].

In order to transmit each bit there is the need for: 13 CW laser sources, 240 TXs and RXs and 720 MRRs. These resources are replicated as many times as the target bit parallelism, and also for the optical clock support. Power metrics of all basic blocks of our architectures are summarized in Table 6.2. The derived static and dynamic power values for electronic and optical components are combined with system-level simulation results (see section 7) to obtain comprehensive metrics under the effect of functional traffic.

## 7 Experimental results

In this section we propose our results in terms of **performance** and **energy consumption**. We followed two fundamental analysis strategies, hereafter

<b>ELECTRONIC DEVICES</b> (// bit parallelism)	<b>STATIC POWER</b> (mWatts)	<b>DYNAMIC ENERGY</b> (fJ/bit)
DC_FIFO_TX_5 //3	0.12	10.65
DC_FIFO_RX_5 //3	0.12	8.54
DC_FIFO_TX_22 //3	0.12	39.00
DC_FIFO_RX_15 //3	0.12	26.50
MUX4x1_ARB //3	0.08	0.36
MUX45x1_ARB //3	0.9	5.09
SERIALIZER //3	0.0475	9.41
DESERIALIZER //3	0.0289	7.74
MESO_SYNCH //3	0.082	8.00
BRUTE_FORCE //3	0.004234	1.4
DC_FIFO_TX_5 //4	0.12	12.72
DC_FIFO_RX_5 //4	0.12	10.2
DC_FIFO_TX_22 //4	0.12	46.41
DC_FIFO_RX_15 //4	0.12	31.65
MUX4x1_ARB //4	0.11	0.49
MUX45x1_ARB //4	0.9	5.09
SERIALIZER //4	0.0417	2.63
DESERIALIZER //4	0.0281	6.12
MESO_SYNCH //4	0.113	11.1
BRUTE_FORCE //4	0.00503	1.66
DEMUX1x3	0.000725	0.92
DEMUX1x15	0.0021	25.21
DEMUX1x4	0.00056	6.72
COUNTER@4bits	0.02964	1.014
TSV	/	2.5
TRANSMITTER (aggressive)	0.025	20
TRANSMITTER (realistic)	0.100	50
RECEIVER (aggressive)	0.050	10
RECEIVER (realistic)	0.150	25
THERMAL TUNING /MRR @20°K	0.020	/
E-SWITCH (3VCs)	17.9	193

Table 6.2: Static and Dynamic Power of Electronic Devices.

referred as: **CMF (Common Modeling Framework)** vs. **AMF (Accurate Modeling Framework)**. Their comparison is very instructive.

The first one reflects common modeling assumptions in the open literature, which lead to an overly optimistic assessment of optical interconnect technology. In particular, network interfaces are typically oversimplified, and end up being abstracted by simple input/output FIFOs of infinite length. Similarly, the blocking effect of the backpressure mechanism is overlooked. As a consequence, the ONoC easily proves much more performance-efficient than the electronic counterpart. Moreover, the lack of a layout analysis in addition to a physical-layer analysis in ONoC design is another important source of optimism in previous evaluations.

In contrast, the key strength of this work (AMF methodology) consists of a careful exploration of E/O and O/E interfaces, accounting for the contributions and effects of every building block: routing, buffering, serialization and deserialization processes, as well as optical transmitters and receivers, clock domain synchronizer, backpressure cost. Last but not least: the propagation of source-synchronous clocks in optical networks and their impact on the overall system power consumption. It will be very interesting to compare the design trade-offs with the two experimental methodologies. They will be hereafter presented with both the conservative and the more aggressive parameters of the optical technology.

## 7.1 Methodology

Experimental results are obtained through GEM5 full-system simulator [76] where we modeled in details both the electronic baseline and the optical architecture described in the previous sections. Modeling included functional behavior, timing accuracy, and energy consumption. Performance and energy are evaluated for the PARSEC 2.1 benchmark suite, a collection of heterogeneous multithreaded applications spanning different emerging application domains [77]. These benchmarks, comprising artificial visions, media processing, 3D and physical animation, search similarity and featuring parallel algorithms, are representative not only of chip multiprocessor workloads, but also of the current and near future usage of high-end embedded devices such as smartphones and tablets [78]. We adopted the medium input-set to have a

significant workload size and maintain a reasonable simulation time (within a few days per benchmark). All benchmarks have run on a Linux 2.6 operating system, which is booted on the simulated architecture, and they are instantiated with a degree of parallelism 16. Benchmarks are modified to enforce core affinity as to avoid non-determinism due to operating system scheduling. As for performance metrics, we considered execution time of the entire parallel region of each benchmark as representative of the end-user perceived performance, and the overall energy consumption of the considered networks. Finally, the shared cache resources are coordinated by a state-of-the-art MOESI coherence protocol.

Cores	16 cores, 1 GHz
L1 caches	16 kB (I) + 16 kB (D), 2-way, 1 cycle hit time
L2 cache	4 MB, 8-way, shrd/dstrb 16x256kB banks, 3/12 cycles tag/tag+data
Directory	MOESI protocol, 16 slices, 3 cycles
eNoC	2D-mesh, 1 GHz, 1 cycle/hop
ONoC	3D, 28 mm length, 16 I/O ports, 10 GHz, 13 wavelengths, 16 waveguides, 430 ps full round
Main memory	4 GB, 300 cycles

Table 6.3: Parameters of the simulated architecture.

Table 6.3 summarizes some architectural details and parameters of the considered reference MPSoC. In the simulator, we implemented the buffering structures of the E/O interface, as explained in the chapter.5. These are needed to correctly support the photonic communications and to precisely take into account the contention and flow-control effects towards a specified destination. We integrated 5-slot and 22-slot FIFOs on the transmitter and receiver sides to model the adopted credit-based approach. The detailed

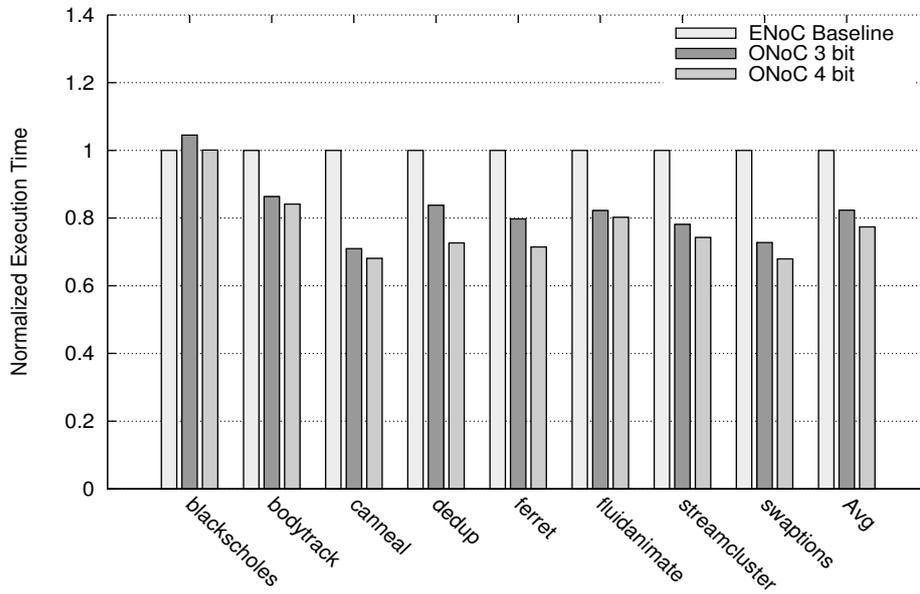


Figure 6.6: Performance comparison of the ONoC with the electronic baseline.

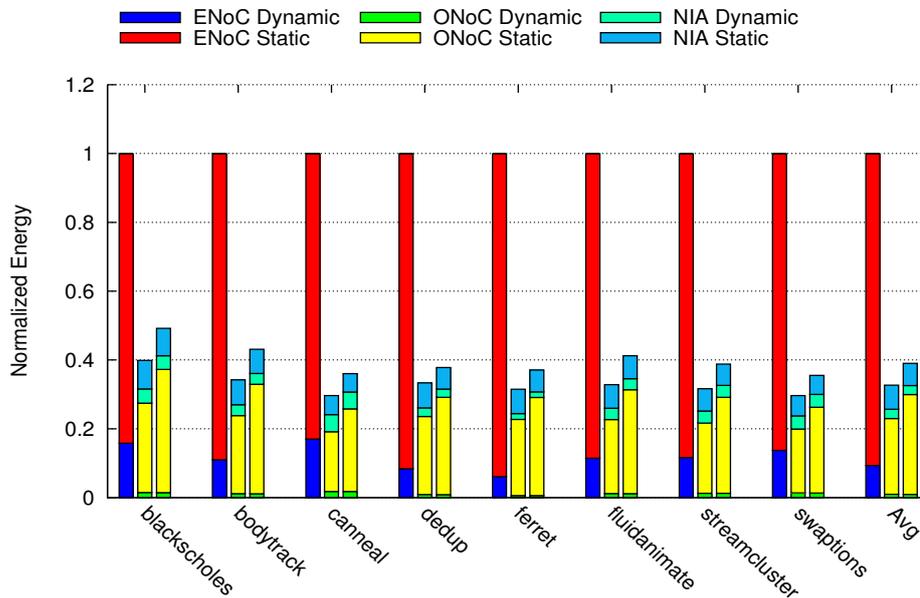


Figure 6.7: Energy comparison of the 3 bit (2nd bars) and 4 bit (3rd bars) ONoC wth respect to ENoC baseline for the common aggressive.

power models of both the electronic and optical networks are integrated, comprising all the details about the network interfaces.

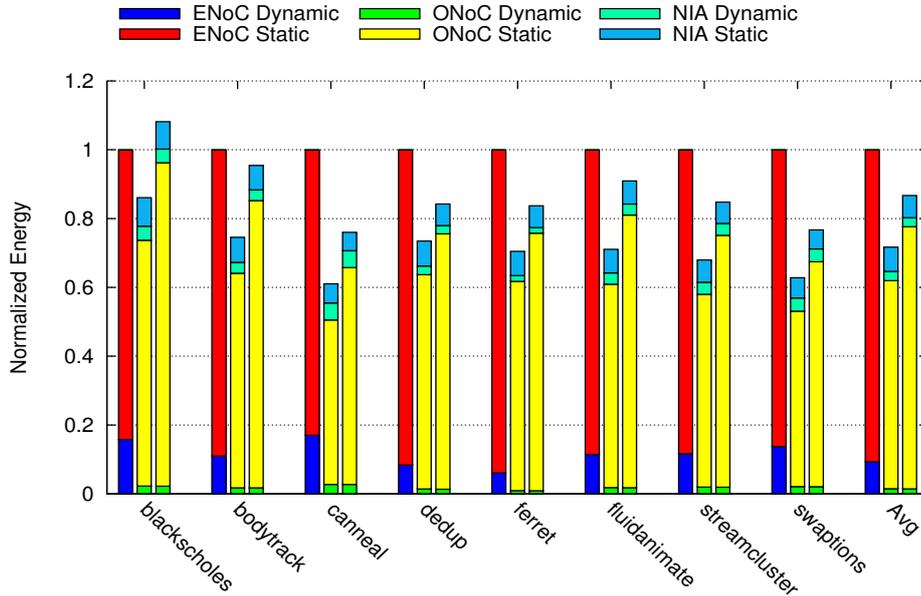


Figure 6.8: Energy comparison of the 3 bit (2nd bars) and 4 bit (3rd bars) ONoC with respect to ENoC baseline for the common realistic.

## 7.2 Result discussion

As shown in Figure 6.6, the optical solution, for both 3 and 4 bit parallelism, is able to deliver performance speedups over the electronic baseline. It achieves up to 23% improvement in case of the larger parallelism (4-bit), with peaks of more than 30% for *canneal* and *swaptions* applications. This speedup is indirectly useful to reduce the overall static energy consumption of the optical network. The 3-bit parallelism scores slightly worse obtaining a 18% performance improvement. We consider ONoCs between 3-bit and 4-bit parallelism, since the bandwidth of optical paths ends up being less than that of electronic counterparts. More than 4 bits are equally not considered since the corresponding static power of the ONoC became unacceptable.

From the energy consumption point of view, Figure 6.7 and Figure 6.8 show the achieved results for the common optimization analysis. In this setup the aggressive and the realistic case show a very different behavior. In the former case, the ONoC saves energy with respect to the electronic baseline by almost 70% on average for the 3-bit case and 60% for the 4-bit one, (see Figure.6.7) by exploiting its reduced static consumption. In the latter case, ONoC also obtains a good energy improvement (avg. of 28% for 3-bit setup and 13% for the 4-bit one) over the electronic baseline demonstrating that,

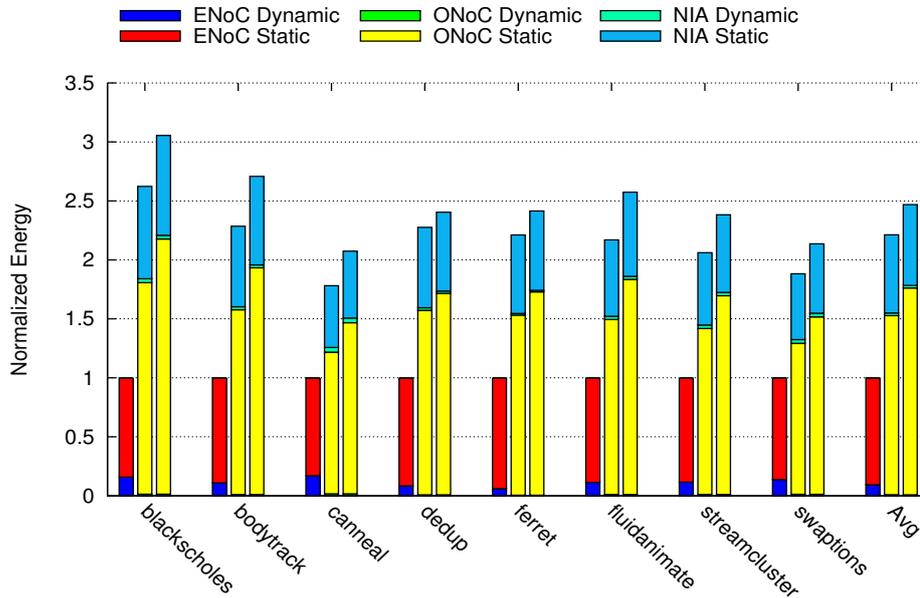


Figure 6.9: Energy comparison of the 3 bit (2nd bars) and 4 bit (3rd bars) ONoC with respect to ENoC baseline for the accurate realistic.

even with the realistic setup, it is still able to get some benefits compared with the ultra-low-power electronic baseline (see Figure.6.8). Although the CMF raises expectations that are not justified in practice, it is already able to point out a realistic effect: in the presence of a coherence traffic, the great dynamic energy savings of the ONoC do not count a lot in the final energy balance, since static power is the dominating factor, and it is mainly associated with the amount of instantiated resources, as well as with technology maturity.

Figure.6.9 and Figure.6.10 show the achieved results for the accurate model. In this case the aggressive setup causes the ONoC to track closer (avg. overhead of 11.6%) the break-even with the very low-power electronic baseline, by exploiting the benefits derived by the execution time speedup of the optical network (see Figure.6.10). The results achieved in the realistic case are instead far from those of the electronic baseline, and this is due to the more relaxed technology used in this setup, and to the more accurate modeling of the optical structures needed. The energy overhead worsens by more than 2x with respect to the electronic energy figures for the 3-bit configuration (see Figure.6.9).

With respect to the AMF, the CMF points out a clear underestimation of

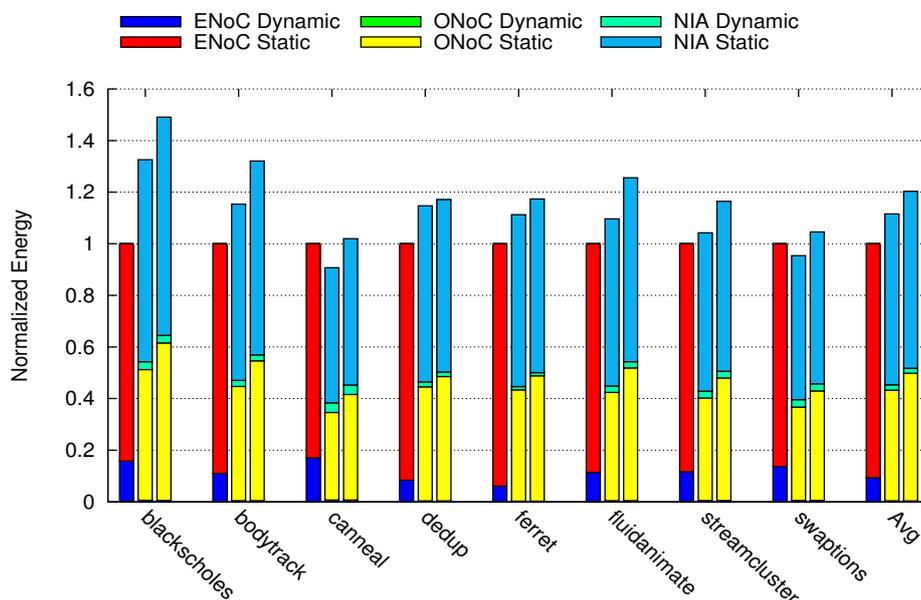


Figure 6.10: Energy comparison of the 3 bit (2nd bars) and 4 bit (3rd bars) ONoC wth respect to ENoC baseline for the accurate aggressive.

static power, which further emphasizes the dynamic power benefits, and especially that the actual complexity and overhead of network interfaces are typically overlooked. Indeed, these results question the common conclusion that ONoC prime time will depend only on the progress of technology maturity. In contrast, such prime time will only come through an in-depth optimization of network interfaces, where the real complexity is hidden under subtle issues such as buffer numbers and sizing, protocol-dependent deadlock, synchronization and flow control. The obtained performance benefits are not able to reduce the static consumption to get results comparable with the very low-power electronic baseline with the nowadays available technology. The good performance results could be however a good starting point in case of future enhancement, giving a hint to better explore and refine the technology applied to these new optical solutions, trying to obtain the maximum advantages from them. Especially, it will be mandatory to explore static power gating solutions, similarly to what happened in the past with electronic circuits.

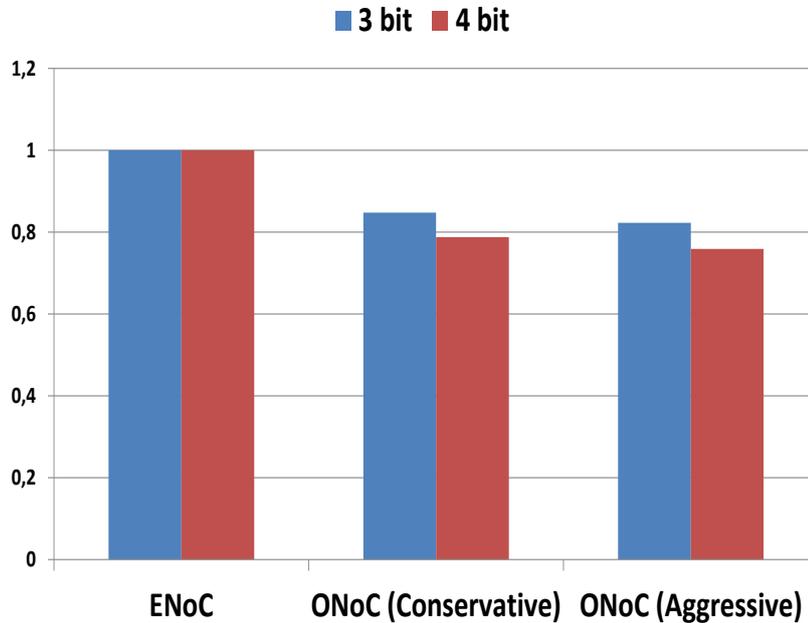


Figure 6.11: Normalized System-Level Energy Comparison.

### 7.3 System-Level Energy and Conclusion

When we extend the focus to the system as a whole, previous results paint a less dismal picture. An interconnect fabric is in fact only a small portion of the total system energy. As discussed in the previous section, the ONoC is capable of speeding up the average execution time (from 18 to 23%, for 3-bit and 4-bit parallelism respectively), meaning that the system as a whole can burn power for less time compared with the electrical counterpart. Figure 6.11, shows normalized the system-level energy contrasting between the two systems, the first one with the ENoC and second one with the ONoC assuming both technologies. In this comparative analysis it is assumed that systems, without the interconnect, consume the same amount of power, (i.e. 15 Watts) which also includes the power spent by the processor elements. As can be seen, ONoC makes the system more energy efficient than the ENoC counterpart regardless the bit parallelism. In particular at 4-bit parallelism, its energy saving results to be **21%** assuming a conservative technology, while with an aggressive optical technology it improves up to **24%**. **This**

---

meas that, when the superior performance properties of optical interconnect technology are properly exploited at the architectural level, certainly, the system will burn power for a lower amount of time. In light of this, an aggressive technology should not necessarily be adopted since energy savings are already there even with a conservative optical technology. Finally, this chapter has included contents that are referred to a cooperative and interdisciplinary work where further details are in [36].



# Chapter 7

## CAD Support for Design and Validation of Optical Networks-on-Chip

### 1 Why an Automatic Place&Route Tool for ONoC Design is needed?

Optical Networks-on-Chip (ONoCs) are considered a promising way of improving power and bandwidth limitations in next generation multi- and many-core integrated systems. Today, most related research acknowledges the key role of the physical layer in assessing ONoC topologies (e.g., insertion loss), but overlooks the placement and routing stage in the design process, hence applying physical design considerations to topology logic schemes. Such a mismatch is fundamentally due to the lack of mature CAD tools for placement and routing of optical NoCs. The objective of this chapter is to bridge this gap. We propose **PROTON**, a fast tool for automatic placement and routing of ONoC topologies, which can support designers in quantifying the degradation of design quality metrics when moving from topology logic schemes to their physical implementation. This gap is especially relevant for Wavelength-Routed ONoCs (WRONoCs), where logic schemes typically make unrealistic assumptions about the placement of initiators and targets. For this reason, we put PROTON to work with the most promising WRONoC

topologies and explore their physical design space given the placement and routing constraints of a 3D stacked system. This chapter addresses a comparative analysis between automatically generated layouts with handcrafted ones for the same topologies and target system, and prove an insertion loss improvement by up to 150x. With PROTON the exploration of the physical design space of ONoC topologies is possible as well as their scalability analysis considering the layout.

## 2 Introduction

Today, for the first time the integration of a fully functional photonic system on a VLSI electronic die can be realistically envisioned. However, despite the arguments in favour of optics for interconnects on the silicon chip and the promising integration route, there is essentially no practical use today. This is in conflict with the vast amount of works in the literature addressing system level redesign around an optical interconnection network and the associated power and performance benefits [32], [31], [29], [65]. In fact, these contributions just foster the optical network-on-chip concept, but are not capable of bridging the gap with an actual interconnect technology of practical relevance. Clearly, further research contributions are needed closer to the physical layer. On one hand, there are still open challenges for the reliable manufacturing and safe runtime behavior of optical components within industrial products. For instance, integrated sources should sustain the working temperatures of a CMOS chip. The problem is also common to passive silicon structures, which in addition suffer from large sensitivity to dimensional variations, calling for trimming or active tuning. Finally, it is worth mentioning the energy improvements that are still required in CMOS compatible receiver circuit design to drastically reduce the energy/bit contribution of the optoelectronic conversion [13],[27].

On the other hand, even assuming industry-standard technology maturity of ONoC building blocks such as laser sources, modulators, detectors and switching elements, the key for the success of optical networks-on-chip consists of a suitable support to exploit the new interconnect technology for system-level design [62], [63]. In this respect, there is today a relevant gap in

terms of design technology, which currently prevents any realistic roadmap for industrial uptake. In this domain, and similarly to digital electronic design in nanoscale technologies, the key concern for the design of an optical interconnection network consists of its predictability. This can be defined as the deviation between the topology logic scheme and its physical implementation in terms of physical parameters such as number of crossings, and the waveguide length. Such parameters are associated with the losses that optical signals experience across the physical paths of the topology. If we define the critical path of an optical network as the physical path with the largest optical loss, then we derive that such critical path determines the minimum power requirement on the laser source, given a specific detector sensitivity. It is therefore of the utmost importance to preserve the predictability of optical losses across design layers to avoid ONoC topologies that map inefficiently to the physical layer. The main source of deviation of real topology layouts from the associated logic schemes consists of an hardly predictable increase of waveguide crossings. There are three main reasons for this. On one hand, logic schemes often make unrealistic assumptions on the positions of initiators and targets on the actual floorplan [52]. Alternatively, some logic schemes are optimized for hardwired positions of the network interfaces, which might not be verified in the system at hand [28]. On the other hand, multilayer photonics is still far from becoming a mainstream solution due to its many fold manufacturing challenges [74] (e.g., sensitivity to process variations in physical gap design). Therefore, single layer photonics might be the reference solution for some time, which increases the risk of additional crossings when routing intricate connectivity patterns. In order to cope with these challenges, designers today have no other choice than placing and routing ONoCs manually, thus basing the insertion loss minimization goal entirely on their intuition and experience. This is due to the lack of automatic placement and routing tools for ONoC design. This chapter takes on this challenge and proposes PROTON, a CAD tool for the automatic placement and routing of optical switching elements and waveguides respectively. Thanks to the flexibility of its objective function, PROTON paves the way for the automatic physical design space exploration of optical NoC topologies. Similar to electronic NoCs [73], automatic placement and routing are need-to-have features for the physical design

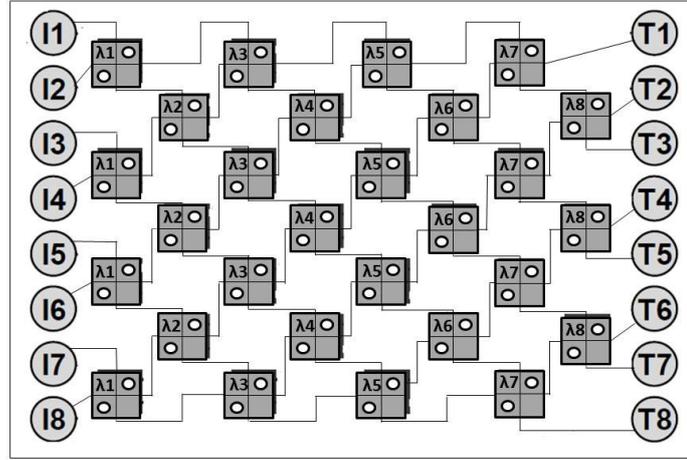
of those topologies that show a clear discrepancy between the logic design and the physical layout. This concern is especially critical for wavelength-routed ONoCs, which fundamentally consist of add-drop optical filters, and therefore end up in multistage interconnection networks that map inefficiently to a 2D surface. Without lack of generality, this chapter will address these topologies, in that they represent the most challenging benchmark for an automatic placement and routing too. The rest of this chapter is structured as follows: Section.3 presents the properties of PROTON starting from of the topology specification format until the PROTON's placement and routing algorithm. Experimental results are discussed in Section.5 before concluding the chapter in 6.

### 3 PROTON's properties

This section goes through the key properties of PROTON.

#### 3.1 Topology Specification Format

In order to process any optical NoC topology, PROTON uses a **Topology Specification Format**, which defines the physical path that is taken by each optical signal at a specific wavelength. Let us consider the 8x8  $\lambda Router$  as a case study. As illustrated in Figure.7.1 each initiator uses 8 distinct wavelengths to reach all 8 destinations, each one following a different routing path. The 8x8  $\lambda Router$  is broken down into those 64 routing paths. As shown in Figure.7.2, each of them connects one Initiator (**I**) and one Target (**T**) by passing several optical filters. In the following, optical filters are called photonic switching elements (PSEs). The topology specification format is being implemented into C++ language to cope with the inherent limitations of a manual description of a given topology logic scheme. In particular, the milestone  $\lambda Router$  topology was completely implemented into C++ language, and the proposed code is able to automatically generate a routing path description. All of informations of the ongoing path such as number of the PSE, its specific order as well as its functionality (cross or drop) are taken into account. This is clearly one of the preliminary steps that should be pursued

Figure 7.1: 8x8  $\lambda$ Router logic scheme.

to develop automatic synthesis tools for optical NoC. Moreover PROTON needs of additional **entry levels** such as dimensions and positions of all optical devices comprising the Initiators and Targets (e.g Hubs and Memory controllers as widely explained in previous chapters), width and height of PSEs as well as the dimensions of the optical layer. By combining topology specification format (also called routing-path information) and entry levels PROTON can generate a valid physical layout of any WRONoC topology, e.g. all PSEs are placed overlap-free inside the die area and are connected by waveguides.

### 3.2 Placement & Routing Algorithm

One optical path consists of the connection between one Initiator (I) and one Target (T). As shown in Figure.7.2 one path consists of several nets (from  $n_0$  up to  $n_k$ ), which are defined as two-pin connections between one Initiator (I1) and one PSE, one Target (T4) and one PSE or two PSEs. Since insertion loss is the most important quality metric in optical NoCs an optimal physical layout of the optical layer minimizes maximum insertion loss (ILmax) over all paths. Thus, the objective function is given as:

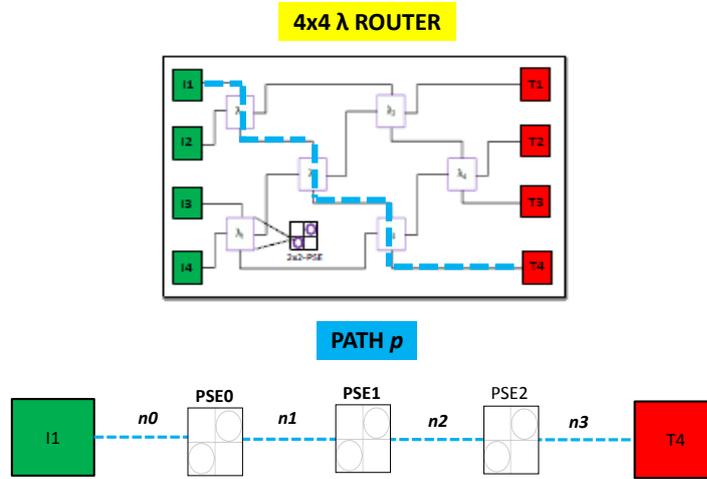


Figure 7.2: An example of optical paths.

$$\text{minimize } \mathbf{IL}_{\max}(\mathbf{x};\mathbf{r}) = \max_{p \in \mathbf{P}} \mathbf{IL}_p(\mathbf{x};\mathbf{r}) \quad (1)$$

with  $\mathbf{P}$ ,  $\mathbf{x}$  and  $\mathbf{r}$  describing the set of all paths, positions of all PSEs and positions of all waveguides respectively. The insertion loss of path  $p$  mainly depends on propagation loss  $\mathbf{pl}_p(\mathbf{x};\mathbf{r})$  and crossing loss  $\mathbf{cl}_p(\mathbf{x};\mathbf{r})$ . By decreasing the crossing loss can result in an increase of propagation loss and vice versa. An example is given in Figure.7.3.

More in details, the PSE0 is connected to PSE1 and PSE2 is connected to PSE3. Figure.7.3 (left) propagation loss is minimized, e.g. the length of waveguides is minimal, and one crossing appears. Figure.7.3 (right) minimizes crossing loss, and crossings are avoided as far as possible. The length of the waveguides increases compared to the routing solution shown in Figure.7.3 (left). Ultimately, minimizing propagation and minimizing crossing loss mostly are contradictory objectives. Our algorithm aims at finding a trade-off by minimizing a weighted sum.

$$\mathbf{IL}_{\max}(\mathbf{x};\mathbf{r}) = \alpha \cdot \mathbf{pl}_p(\mathbf{x};\mathbf{r}) + \beta \cdot \mathbf{cl}_p(\mathbf{x};\mathbf{r}) \quad (2)$$

with  $\alpha, \beta$  real numbers and  $\alpha + \beta = 1$ .

Due to the complexity of the placement and routing problem our algorithm

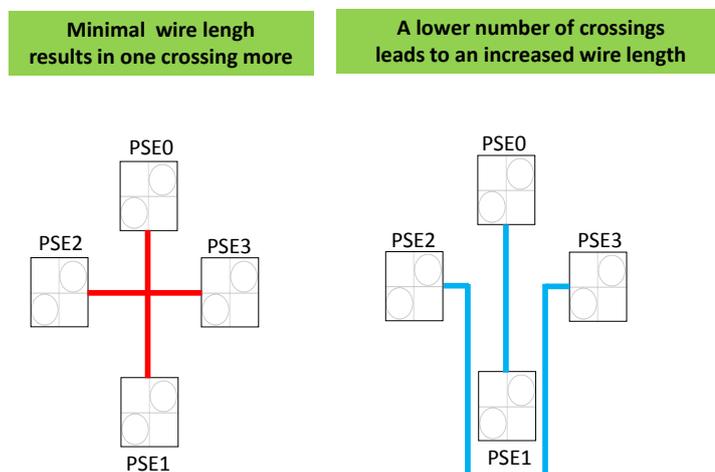


Figure 7.3: Propagation and Crossing Loss are tightly interrelated.

is structured into two steps: During the first step suitable positions for PSEs are found considering the distances between two PSEs and the (approximated) number of waveguide crossings. After this the waveguides are routed by minimizing the length of the waveguides and/or the number of crossings. Additional details about PROTON's placement and routing algorithm are reported in [39].

## 4 Maximum Insertion Loss

After routing PROTON counts the number of crossings as well as drops and determines the length of waveguides for each path. By using the loss parameters reported in chapter.4 it is possible to derive  $IL_{max}$  (sum of all loss contributions which are closely related to their amount ) across all paths of the ONoC under test. By improving the optical technology in the future, only loss parameters have to be adapted. Consequently, PROTON results to be a flexible and technology independent tool for the design of ONoCs.

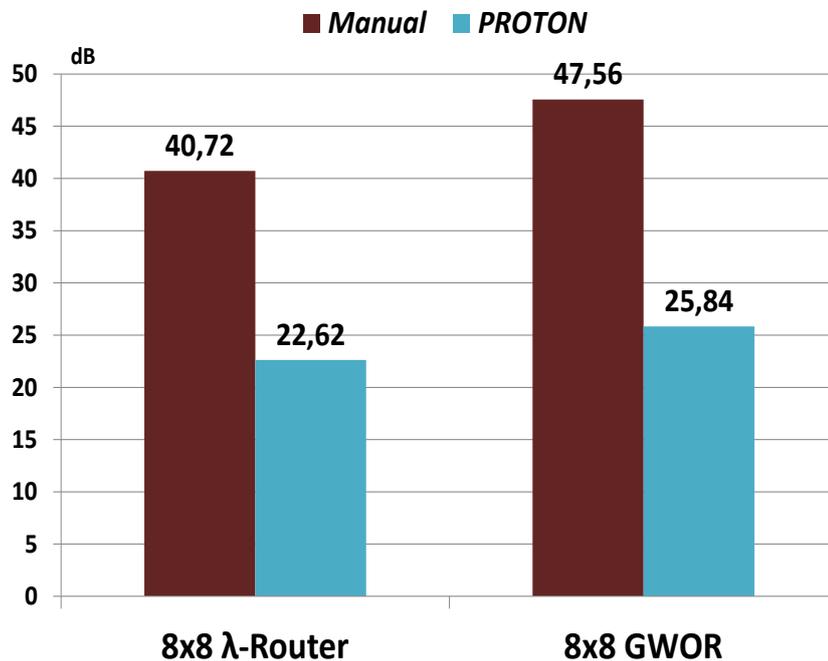


Figure 7.4: ILmax (dB) Contrasting: Manual vs. PROTON.

## 5 PROTON at work

The algorithms are implemented in C++ and all experiments are executed on an Intel Core 2 Quad CPU with 8GB RAM running at 2.33GHz. For solving the optimization problem we use the IPOPT library [75] which is one of the leading libraries in nonlinear optimization. This section shows: the comparison between PROTON and a manually created layout. In section.?? a study about the calibration of PROTON and the best topology selection is given. In section.5.3 the scalability of PROTON is proven.

### 5.1 Manual Design vs. PROTON

In this section, as a case study, we implement the 8x8  $\lambda$ Router and the 8x8 GWOR using two relevant methodologies: Manual design (as widely addressed in chapter 4) vs. CAD tool. As insertion loss is the most important physical metric that must be quantified to determine the laser power that

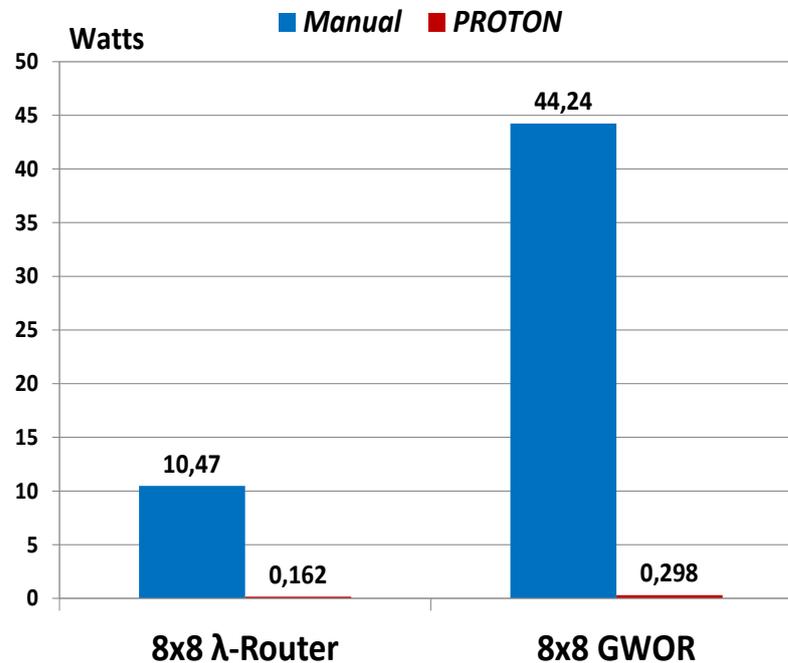


Figure 7.5: Laser Power (Watts) Contrasting: Manual vs. PROTON.

guarantees a predefined bit error rate at receivers, we assess the worst case  $IL_{max}$  for each topology. Our study therefore assumes the loss parameters mentioned in chapter 4. Crossing loss and drop loss were obtained by 2D Finite-Difference-Time-Domain (FDTD) method. Since every path meets at most one drop and drop loss is very small compared to crossing- and propagation loss it is considered to be negligible. Similarly, we do not consider bending loss in our calculations. Moreover, in both methods, we optimize the number of crossings so that the only one objective function being optimized is the crossing loss. We also optimize every crossing with the standard elliptical taper [60]. We also take for granted the same architecture, hypothesis, and assumptions already followed in the chapter 4. Figure.7.4 shows the insertion loss deviations between manual layout and the automatically generated for both topologies.

As can be seen manual design and CAD tool feature the same trend. In particular,  $\lambda Router$  provides lower insertion loss than GWOR in both cases,

40.73 dB vs. 47.58 dB in the manual layout, while 22.62 dB vs. 25.84 dB in the CAD tool. This also optimizes the insertion loss by 65x over the manual layout for  $\lambda Router$  topology, while by 150x for the GWOR one. The benefit of this effect can be seen when quantifying the laser power requirements. In fact, once the maximum insertion loss across all wavelengths and optical paths is obtained and the detector sensitivity (S) is known, we finally evaluate the lower limit of optical laser power (P) to reliably detect the corresponding photonic signal at the receiver end. In our evaluation we make the practical assumption that such a worst case ILmax dictates the power needed by all laser sources. For laser sources we assume a laser efficiency (PLE) of 20% and a coupling laser-link (PCW) of 90%, while for detectors we consider a sensitivity (S) of -17dBm with a BER= $10^{-12}$ . As shown in Figure.7.5 in both cases  $\lambda Router$  is clearly most efficient than GWOR due to the lower insertion loss. We note 10.49 Watts vs. 44.44 Watts in the manual layout, while 0.162 Watts vs. 0.298 Watts in the CAD tool. Thanks to the lower insertion loss that CAD tool exhibits with respect to the manual layout, the lasers power requirements are optimized by 98.5% in the  $\lambda Router$  topology and by 99.3% in the GWOR one. Because the topology is too complex to determine an optimal physical layout manually we obtain much better results using CAD tools. Furthermore PROTON needed only a few seconds while creating the manual design took approximately one week per topology.

## 5.2 Best Topology Selection

PROTON minimizes the weighted sum of propagation loss and crossing loss as can be seen in equation (2). Thus, the tool can be calibrated by choosing different values for  $\alpha$  and  $\beta$ . Table I shows the results for 8x8  $\lambda Router$ , 8x8 GWOR when varying those weights. In this study, also the Standard crossbar is considered. All topologies are placed on a die area of 8mm x 8mm. Elliptical tapers are assumed for all crossings. The first two columns contain the varied weights  $\alpha$  and  $\beta$ . The column ILmax gives the maximum insertion loss. The number of crossings and the waveguide length in  $\mu m$  of path p with  $IL_p = IL_{max}$  can be seen in column CR and WL respectively. At last CPU time in seconds is presented. As can be seen in the table, maximum insertion loss is high when the number of crossings is high. Thus, insertion loss

Weights		8x8 $\lambda$ -Router				8x8 GWOR				8x8 Standard-Crossbar			
$\alpha$	$\beta$	$il_{max}$	CR	WL	CPU	$il_{max}$	CR	WL	CPU	$il_{max}$	CR	WL	CPU
1.0	0.0	39.8	73	12402	77	45.2	81	20178	47	40.3	70	25821	351
0.9	0.1	18.1	32	9324	76	23.4	42	10161	50	25.0	45	11052	352
0.8	0.2	25.7	46	11682	79	24.5	44	10476	50	32.4	55	25362	352
0.7	0.3	24.5	44	10413	76	25.6	46	11088	50	25.0	41	24516	359
0.6	0.4	15.5	26	13392	76	25.1	44	14589	53	24.8	41	23499	356
0.5	0.5	23.8	42	13203	76	29.8	54	11241	51	25.6	46	10782	382
0.4	0.6	17.2	30	10539	76	28.8	52	11889	50	27.7	48	18513	352
0.3	0.7	25.0	44	13824	77	26.7	48	11610	51	28.4	48	22842	353
0.2	0.8	22.1	39	12366	77	24.7	44	12258	52	30.6	53	20007	354
0.1	0.9	22.2	39	12681	78	22.0	38	14886	50	24.7	42	19269	351
0.0	1.0	22.6	41	8676	78	25.8	43	23094	51	27.2	45	25578	353

Figure 7.6: Table I: Results for variation of propagation and crossing weights.

mainly depends on the number of crossings. Comparing all three topologies  $\lambda Router$  is the one with lowest ILmax, followed by GWOR and Standard-Crossbar. Hence, 8x8  $\lambda Router$  is the optimal topology to be chosen in terms of maximum insertion loss. PROTON acts very fast on all three topologies with a maximum runtime of 6.3 min. CPU time mainly depends on the number of PSEs and nets. Since number of nets for 8x8 GWOR is lower than for 8x8  $\lambda Router$  and 8x8 Standard-Crossbar PROTON needs lowest run time for 8x8 GWOR.

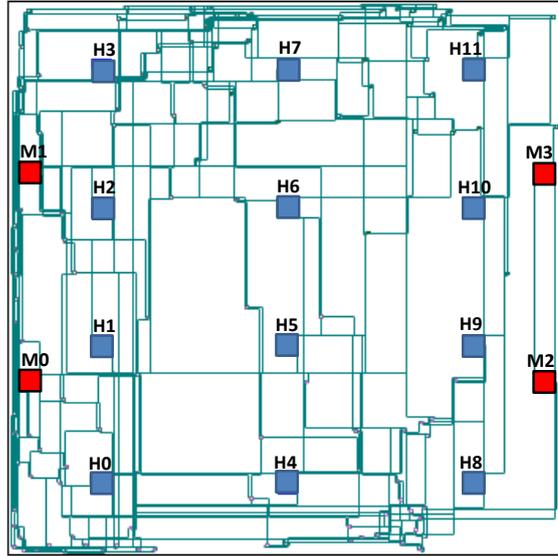


Figure 7.7: 16x16  $\lambda Router$  under scalability assumptions.

### 5.3 Scalability

In this section we characterize the impact of system scale and we show that PROTON can handle large topologies. As a case study we test the 16x16  $\lambda Router$  as the best topology of our exploration. We sketch a future generation of our target system, which now consists of 96 cores in the tiled-based electronic layer, while preserving 8 cores for each cluster. For this purpose 12 gateways (or 12 hubs) are needed in the optical plane to directly connect to the corresponding cluster of the electronic plane. Moreover thanks to the benefits of photonic integration deeper into the DRAM DIMM [30] we kept the same number of memory controllers. Note the new die size of 12 mm x 16 mm. Results are shown in Figure.7.8. The objective function includes all kinds of loss optimization: propagation loss, crossing loss and finally their combination. The first three bars show the ILmax across different optimizations, while second and third ones show the number of crossings and the waveguide length of path  $p$  with  $IL_p = IL_{max}$  respectively. CPU time can be seen in the last three bars. All lengths are measured in cm and the

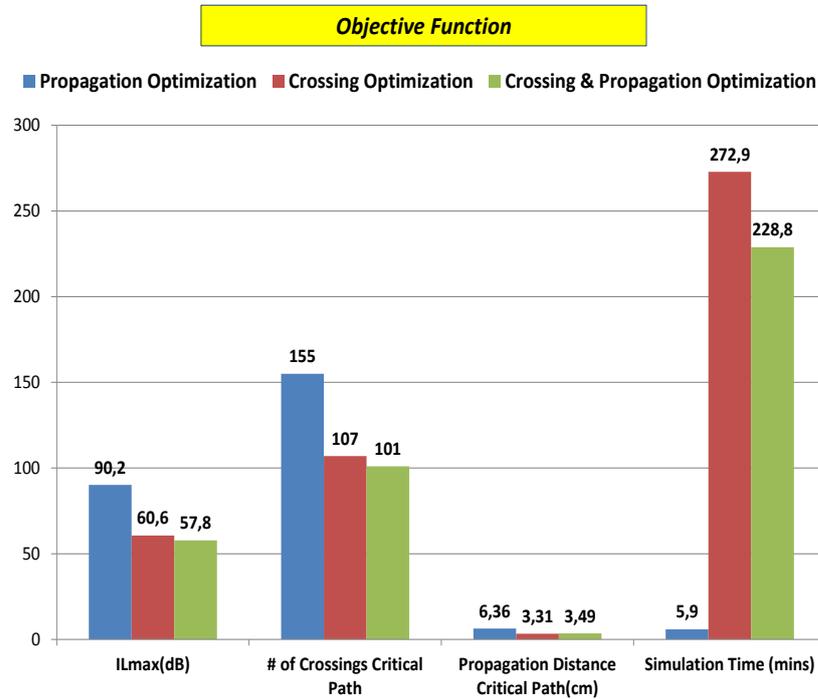


Figure 7.8: Results under scalability assumptions.

times in minutes. Although the number of paths is 5.2x higher compared to 8x8  $\lambda Router$  maximum insertion loss increases less than 2.6x. Figure 7.7 shows the layout of 16x16  $\lambda Router$  when minimizing propagation or crossing loss. The blue rectangles are Hubs while the red ones are the Memory controllers. Optical filters are small rectangles that are connected by waveguides illustrated with green lines. PROTON places and routes fast with a worst case execution time of at most 272.9 minutes. For larger sizes such topologies would become unusable due to an insertion loss exceeding 91dB on the critical path.

## 6 Conclusion

This chapter introduced the first automatic tool for placement and routing of optical NoCs. **PROTON** iteratively places all optical filters overlap-free inside the chip area and routes the waveguides both while minimizing prop-

agation and/or crossing loss. Compared to a handcrafted layout PROTON reduces maximum insertion loss up to 150x while standard topologies are placed and routed in less than 7 minutes.

Therefore, PROTON minimizes the design predictability gap even more than manual design, thus resulting the right methodology for the place&route of filter-based topologies. However, for the ring-based counterpart the manual design still remains the affordable technique.

In the future, PROTON will be improved by post-layout optimization, (e.g. by rotating optical filters to minimize additional crossings), and definition of side fences to avoid undesired crossings at the chip boundaries. Such physical effects are mainly due to the unavoidable intersection between on-chip waveguides and in-out connections needed to enable chip-to-chip communications. Furthermore, the implementation of multimode interference (MMI) tapers will be also considered at every crossing as well as the use of network partitioning. Finally, this chapter has included contents that are referred to a cooperative and interdisciplinary work where further details are in [\[39\]](#).

# Chapter 8

## Network-Level Simulation Frameworks for Optical Networks-on-Chip

### 1 Abstract

This chapter presents a bottom up-abstraction procedure based on the design-flow FDTD + SystemC suitable for the modeling of optical networks-on-chip. In this procedure, a complex network is decomposed into elementary switching elements whose input-output behavior is described by means of scattering parameters models. The parameters of each elementary block are then determined through 2D-FDTD simulation and the resulting analytical models are exported within functional blocks in SystemC environment. The inherent modularity and scalability of the s-matrix formalism is preserved inside SystemC, thus allowing the incremental composition and successive characterization of complex topologies, typically out-of-reach for full-vectorial electromagnetic simulators. The consistency of the outlined approach is verified, in the first instance, by performing a SystemC analysis of a four input and four output ports switch, and making a comparison with the results of 2D-FDTD simulations of the same device. Finally, a further complex network encompassing 160 micro resonators is investigated, the losses over each routing path are calculated and the minimum amount of power needed to guarantee an assigned BER is determined. This chapter lays the basis for

an automatic Technology-Aware Network-Level simulation framework, capable of assembling complex optical switching fabrics, while at the same time assessing the practical feasibility and effectiveness at physical/technological level.

## 2 Background & Motivations

Chip Multicore Architectures currently represent the state of the art in the design of high performance Very Large Scale Integration (VLSI) systems. In accordance with this architectural paradigm, several processing units are physically realized on the same silicon die, and share the execution of the instructions with an high degree of parallelism. For the next generation of digital systems, the International Technology Roadmap for Semiconductors (ITRS) expects the integration on the same substrate of hundreds of computational cores [44]. To assure a fast and reliable communication over a network with such a complexity, a traditional bus-based solution is no more conceivable. From a technological point of view, a straightforward and alternative way to implement the communication among the cores is to create, at the chip level, a network-based link system (Network-on-Chip, NoC) [45]. However in a NoC implemented with electrical links (Electronic Networks on Chip-ENoC), as the number of interconnected cores increases, the constraints in terms of power dissipation and required bandwidth grow exponentially, thus imposing soon unrealistic conditions for practical achievements. To overcome these problems, the realization of optical-based interconnections among cores (Optical Network on Chip-ONoC) appears as a promising solution, both for power consumption and allowable bandwidth, as previously mentioned in chapter.1. By following the last trends of 3D integration in digital systems, it is possible to conceive a vertical stack with a top layer reserved to the optical communication network, superposed to silicon layers incorporating memory and processing units. As discussed in the previous chapters of this thesis, a complex ONoC can be considered as the composition of several elementary building-blocks named Photonic Switching Elements (PSEs). For the design and the optimization of such a complex optical networks, the availability of efficient and reliable tools is fundamental. These tools should

combine the flexibility and the versatility needed to analyze the high level architecture layers, with the capability to represent accurately the basic communication parameters such as, for example, attenuation and bandwidth, which are strictly linked with the physics of these devices. Indeed, taking into account these parameters also at higher level of abstraction (by means of what we identify here as technological-annotation) is fundamental to explore how constraints imposed by the technological aspects can affect the implementation of an optical interconnection network for chip-level integrated systems.

Among the available tools, PhoenixSim[46] is certainly a relevant simulation environment structured in OMNET++. This tool is capable of assessing the performance of multi-processor hybrid systems, integrating electronic and optical networks on the same platform. The key feature of PhoenixSim lies in the modeling parameters, such as: propagation delay, insertion-loss, occupation area on the chip and energy consumption. All these figures of merit are obtained from hardwired pre-characterized values for given wavelengths. For this purpose, PhoenixSim does not use any analytical model for insertion loss analysis.

Although with PhoenixSim it is possible to explore optical on chip networks using a physical layer analysis, it lacks of compliance with industry-standard hardware modeling languages and methodologies. For this reason, PhoenixSim was recently augmented with a SystemC wrapper. On the other hand this procedure is onerous in terms of computation time. In contrast with PhoenixSim, we aim at modeling optical NoCs leveraging on a plain SystemC modeling style. SystemC is an open-source System Level Design Language based on C++, which extends the modeling capabilities of traditional hardware description languages (HDLs) to higher levels of abstraction. SystemC is particularly suitable for the analysis of electronic NoCs and, with some appropriate extensions, can be used also for optical NoCs.

The first example of modeling framework in SystemC is presented in [47]. Here, separate channels are used to model wavelength and power information of optical signals. In [48], a new SystemC class is created to manage analog signals transmitted between modules, which access communication channels through the typical constructs of abstract simulation (e.g. fifo constructs of

untimed functional simulation). Above all, the key issue of technology awareness is solved by invoking the Matlab Symbolic Toolbox for elaborating S-Matrices. In particular, the computation time of all the S-Matrices of each slice of a typical wavelength routed 4x4 optical crossbar is about 30 seconds on a 2.4GHz Pentium 4. With respect to this work, the presented simulation framework aims to reuse the existing port-interface-channel constructs of SystemC, thus making the top-level view of an optical NoC look like the same of an electronic NoC. The difference lies just in module implementations and in the data types exchanged through the pre-defined SystemC channels. Moreover, we leverage on the RTL (Register-Transfer-Level) modeling style for the sake of accuracy. We do not build a new SystemC class to manage analog signals in the network but, in contrast, we exploit the user-defined data type and then we describe the optical information on three different fields such as: logic value, wavelength and signal amplitude. The key challenge of our SystemC modeling framework lies in the integrating technology-annotations in the abstract model to preserve a valuable degree of technology awareness while limiting repercussion in simulation time. Unlike [48], our methodology provides validation results of developed analytical models with respect to FDTD simulations. By doing so, in the first instance, the photonic elements composing the optical links are described at the phenomenological level through a set of analytical models of black-box type, whose parameters are determined through measurements or simulations. Basic optical switching components are then modeled in SystemC through a module that embeds both the functional behavior of the component as well as its non-functional information (i.e.technology-annotations). This allows the exploration of different topologies for the ONoCs, without losing the awareness of the fundamental physical constraints imposed by the optical devices used to compose the optical layer of the network. The basics of the proposed approach are described in next sections; then, this technique is applied to analyze a 4x4 Switch. Comparisons of SystemC results with the ones obtained through electromagnetic simulation of the overall device by Finite Difference in the Time Domain (FDTD) [49] are used to validate the implemented model. A more complex structure (the 4x4 Square Root topology described in chapter.2), realized by compositions of these 4x4 Switches, is therefore investigated and

its insertion-loss and the optical power required by the laser source to meet a fixed detector sensitivity are quantified.

### 3 Technology-Aware SystemC Simulation for Optical Networks-on-Chip

As previously stated, the architectural design of an Optical Network-on-Chip must be performed with a simulation tool able to meet multiple requirements such as efficient network-level simulation, support for technology annotations in compliance with the industrial standards for the design of electronic parts. The SystemC environment offers the necessary features to become the of the art” simulation tool in the system level design of Optical Networks-on-Chip. More specifically the key reasons that make SystemC the most effective choice are:

- (1) SystemC is an Object Oriented C++ class library, that offers a high level of modularity and flexibility. In particular, the base constructs can be custom-tailored to fulfill specific modelling requirements.
- (2) The communication semantics inside SystemC are based on a very flexible set of interface method calls. Consequently, the peculiar and distinctive features of optical links can be captured by leveraging on the preexisting communication constructs of SystemC.
- (3) SystemC can easily span over a wide range of abstractions layers, from Register-Transfer-Level (RTL) up to Untimed Functional (UF), thus allowing both high level of accuracy or reduction of the computation time. Therefore, SystemC may serve as an unified description language able to overcome the limitations of co-simulation approaches. It can provide local and global optimizations, allowing an easier exploration of the whole design space of ONoCs. The bottom-up abstraction procedure leading from the physical models up to their corresponding SystemC modules, can be summarized in the following steps:

- ( $\alpha$ ) Description of the input-output relations of each functional component of the optical network (waveguides, bendings, ring resonators) in analytical

form, by means of the scattering parameters formalism.

( $\beta$ ) Electromagnetic simulation (or experimental characterization) of the elementary functional components, to derive the input-output responses to be reproduced by the analytical models.

( $\gamma$ ) Back-annotation of the analytical models in SystemC, and validation with respect to the numerical or experimental data across the entire optical spectrum.

( $\delta$ ) Modular composition and modelling, inside the SystemC environment, of higher-order routing structures.

( $\eta$ ) Insertion-loss assessment of optical network topologies, and then evaluation of the minimum optical power that the laser sources should provide to enable correct detection of the optical data stream at the photodetectors.

To demonstrate the effectiveness of the proposed approach, a complex 4x4 Square Root topology will be investigated. The analytical model of this network will be obtained by composition of the scattering matrices (s-matrices) of the different elementary building blocks, whose parameters will be deduced through FDTD simulations. This procedure is detailed in next sections.

## 4 S-parameters modelling of a 1x2 PSE

As mentioned above, the most simple PSE is a structure with one input and two possible output ports (1x2 PSE). Without lack of generality, we refer to microring-based 1x2 PSE. In the case of ring-based ONoCs, the 1x2 PSE can be implemented via a microring resonator cascaded to a crossing between two orthogonal waveguides (see Figure.8.1). The device is active, if the resonances of the microring can be dynamically adjusted via some thermal or charge-injection-based effects; otherwise the device is passive, and the microring resonances are fixed and defined a-priori. In the passive configuration, however, a slow thermal mechanism can be envisaged for fine tuning of the resonances in the final set-up of the device.

A microring resonator and a crossing between two waveguides (network  $\alpha$  and  $\beta$  in Figure.8.1) can be both represented as four-ports devices; consequently, the relations between the input and output signals at their ports can

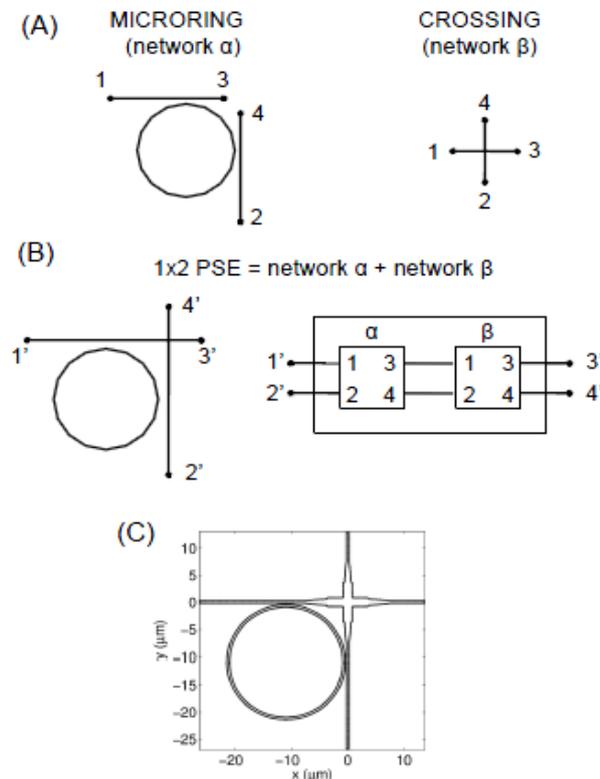


Figure 8.1: (A-left) Sketch of a microring resonator with orthogonal access waveguides. (A-right) Sketch of an orthogonal planar crossing. (B) The 1x2 PSE can be considered as the cascading of the microring the crossing. If the wavelength of the optical carrier signal is resonant with the microring, the data stream on the Port 1' is routed toward Port 2'; if the optical signal is not resonant, the data stream is routed to Port 3'. Dually, a signal coming from Port 4' is routed to Port 3' if resonant, while continuing to Port 2' if out of resonance. The device is bidirectional, i.e. each of the ports can act as input or output, but not simultaneously. (C) Representation of the PSE as it appears when discretized inside the 2D-FDTD code; in this case the transmission at the intersection between the waveguides is optimized using a MMI-based crossing.

be modeled by means of 4x4 scattering matrices, whose coefficients depend on a finite set of parameters (optical lengths, coupling coefficients, transmission efficiency). The resulting s-matrices are symmetrical because the networks

under consideration are reciprocal. Without lack of generality, the gaps between the two waveguides and the ring are supposed equal, and represented with the same power coupling coefficient  $K$ . Moreover, microring and waveguides are assumed dispersionless, thus the effective index and the group index coincide. The propagation losses are also neglected. Dispersion effects, losses, or a coupling asymmetry between the waveguides and the ring can be easily taken into account in the model [50].

Details about s-matrices along with the analytical model are well described in [40]. Simulation approach has been used, and the 1x2 PSE has been characterized through 2D-FDTD. A wide band excitation signal in the wavelength range of 1500nm-1600nm has been applied to Port 1' (see Figure8.1) to evaluate the on-resonance (1'2') and off-resonance (1'3') spectral responses of the device.

Figures.8.2 compares the results of the FDTD simulation (solid lines) with the ones of the s-parameters based analytical model (dashed lines), once applied the optimized fitting procedure. The blue lines refer to the path between the Port 1' and 3' (Through), whereas red lines refer to the path between Port 1' and 2' (Drop). The physical and geometrical parameters of the tested PSE are: internal radius of the ring  $R = 10 \text{ um}$ , ring and waveguides width  $w = 450\text{nm}$ , effective index of ring and waveguides  $n_{\text{eff}} = 2,3561$ , physical gap  $g = 300\text{nm}$  corresponding to a coupling coefficient  $K = 0,0838$ . The MMI-based crossing considered in this PSE provides a transmission efficiency  $\eta = 0,975$ .

## 5 SystemC Modelling of a 4x4 Optical Switch

The integration inside SystemC of the s-parameters model for the elementary 1x2 PSE allows the modular composition and characterization of higher order routing structures. The first composite device under test, whose scheme is illustrated in Figure.8.3, is an optical switch with four input and four output ports, located along the cardinal points: North, East, South and West. Such a structure is realized with eight elementary 1x2 PSEs, arranged in a matrix layout. This switch contains a limited number of waveguide crossings which make propagation losses to be within an acceptable level for a practical

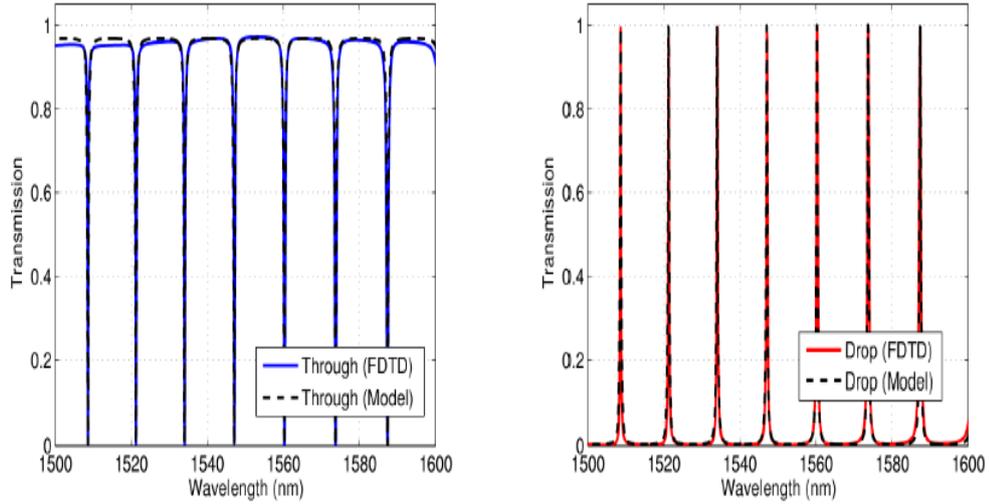


Figure 8.2: Comparisons between the spectral response of the Through (Drop) path [Port 1' to Port 3'] ([Port 1' to Port 2']) of the passive 1x2 PSE, calculated via 2D-FDTD and by means of the s-matrix mode

realization. In order to verify the correctness of the proposed compositional approach, the 4x4 Optical Switch (4x4 SW) has been first simulated with 2D-FDTD. Afterwards, the spectral response obtained from the numerical data has been compared with the one calculated with the SystemC analysis.

By considering the 4x4 SW realized with eight 1x2 PSEs with ring radius of  $10\mu\text{m}$ , the whole footprint of the structure is about  $140\times 140$ . Consequently, the FDTD simulation takes more than 72 hours on a parallel cluster of 10 processors. Due to the burden of the computational domain, the 4x4 SW is, on the available cluster, the upper bound for an FDTD electromagnetic simulation, and represents the top benchmark for testing the reliability of the compositional abstraction procedure. On the contrary, the SystemC modeling

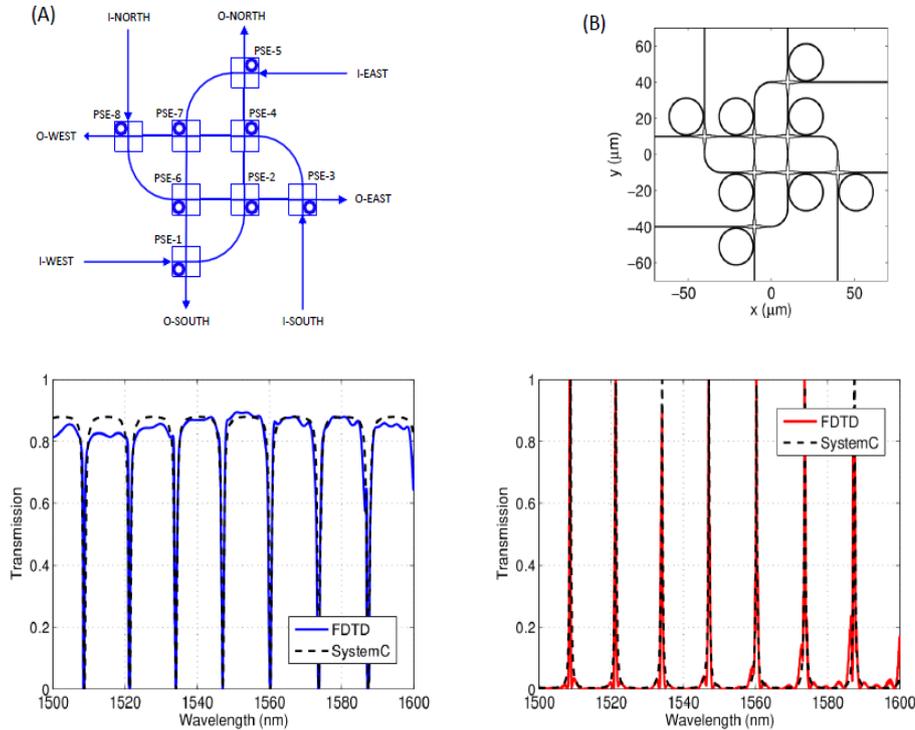


Figure 8.3: Topological scheme of the 4x4 Optical Switch showing the interconnections between the eighth PSEs. (B) FDTD representation of the real device. Down, Simulated (solid line) and modeled (dashed line) transmission curves for the paths linking the I-WEST port with the O-NORTH port and I-WEST port with the O-SOUTH port of the 4x4 Optical Switch.

requires only 0.001 seconds per wavelength to perform the analysis of a single considered optical path.

As illustrated in the panel (A) of Figure.8.3, each 1x2 PSE is connected with other ones by means of specific physical links. However, depending on the wavelength of the signal respect to the resonance of each router, different logical paths must be considered for each connection between the input and the output ports. To describe the overall behavior of the network, SystemC should then model all these logical paths, taking into account all their relevant features. To enable this functionality, we utilize a predefined standard communication channel of SystemC (i.e., sc-signal) with a new user-defined

data type. This allows to replicate the relevant features of an optical link such as logic value (transmission of logical value 0 or 1), optical wavelength and signal amplitude. The optical wavelength is used by the router model to implement the routing functionality. The signal amplitude, on the contrary, is considered to take into account the technological awareness (such as insertion-loss on the direct path and crosstalk determined by spurious power addressed to other ports). To allow the computation of the optical power budget of each link, back annotation from FDTD to SystemC of losses in waveguides, bending and crossing is therefore mandatory.

The leftmost lower side of Figure.8.3 shows the transmission characteristic of the 4x4 Optical Switch, when a signal is injected in the I-WEST port and collected at the O-NORTH port (Through Path). The PSEs involved in this communication are PSE-1, PSE-2, PSE-4 and PSE-5; the data stream is routed through this path if the wavelength of the carrier is not in resonance with the microrings. The solid blue line refers to the result of the FDTD simulation, while the black dotted line is the SystemC modelling. As illustrated by the figure, the SystemC approach fits perfectly the resonances and the global level of losses for the considered path. In a similar way, the transmission characteristic between the I-WEST port and the O-SOUTH port (Drop Path) is shown at the rightmost lower side of Figure.8.3.

Here the signal is routed for wavelengths corresponding to the resonances of PSE-1. It must be noted that, since this topology is symmetrical under step rotations of 90 degrees, the transmission curves are the same injecting the signal from all other inputs (South, East, North). Thanks to the accuracy of the technology-annotations in SystemC, our framework gains control over both the resonant wavelengths of switching components and the signal amplitudes in optical paths. As a consequence, the comparative analysis of wavelength-routed vs. space-routed optical NoCs, or the assessment of Signal to Noise Ratio (SNR) at optical receivers become feasible. In fact, such tasks require the knowledge of both network behavior and of key implementation details and physical insights. The accuracy of the SystemC-based modeling has been further analytically tested by evaluating two different quantities: the average error on the entire spectrum of the signal (SE) and the error on the peaks (PE) (details are in [40]). The former (SE) measures the Mean-

Squared-Error between the FDTD and SystemC spectral responses, and has been calculated across the wavelength range from 1500nm to 1600nm. The mismatch parameter SE achieves 2.45% when is calculated in the case of a Drop Path (for example from I-WEST to O-SOUTH), whereas in the case of a Through Path (as the one from I-WEST to O-NORTH) the SE value is 1.05%. For what concerns PE on the optical spectrum we obtain 0.0063%, stable error for all ports has been calculated. Finally, PE (on the peak values) is around 3,451%. This demonstrates the accuracy of the approach based on FDTD + SystemC for the simulation of complex ONoCs. To further assess the performance of an ONoC, it is fundamental to evaluate the insertion-loss on the critical path (ILmax), which is the path between the input and the output presenting the maximum value of losses. This parameter has been calculated in two different cases, with standard elliptical-tapered crossings and with the MMI-based crossings, in order to evaluate the impact of the technological choice in the realization of the crossings inside the network. In the first case ILmax is around 2dB, while it is reduced to 0.56dB in the second one, thus confirming the better properties of the MMI-based approach to the design of crossings. It is worth noting that the selection between the two possible optimized solutions could also depend on the available space in the physical region around the crossing.

## 6 SystemC Modelling of a 4x4 Square Root Topology

One of the most important examples of optical network topologies proposed in the literature so far, is represented by the 4x4 Square Root as previously described in chapter 2. As illustrated in Figure.8.4, this interconnection network embeds 16 gateways. Each Gateway ( $G_i$ ) can work as the initiator or the target of the communication, and may send and reach optical data from the others. Therefore, a parallel communication is possible. This optical architecture is constructed recursively starting from a 2x2 quad, which in turn consists of four 4x4 switches structured in a 2x2 mesh. In this specific case, the four internal rings of each 4x4 switch have different radius with respect to

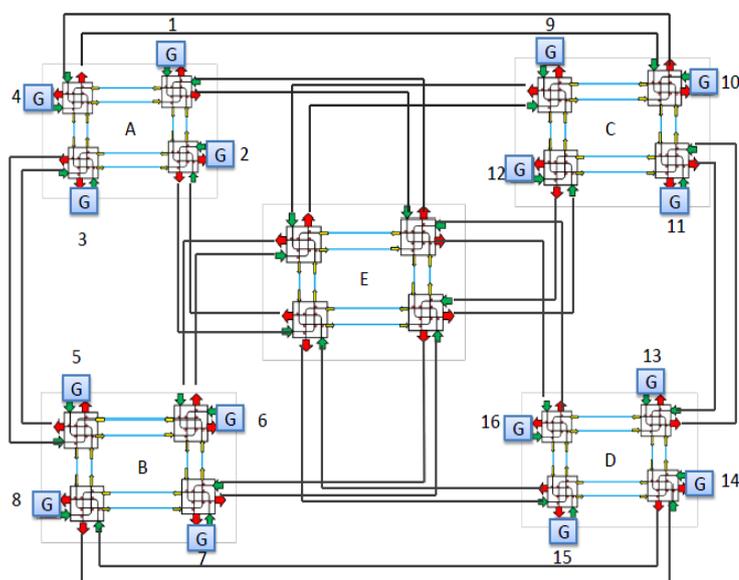


Figure 8.4: 4x4 Square Root Topology.

the four external ones. This choice produces two interleaved resonance spectra, allowing the routing from any input of the 4x4 switch to three possible outputs (i.e. routing from West to South if the signal is resonant with the four external rings, routing from West to East if the signal is resonant with the four internal ones, and from West to North if the signal is not resonant with anyone). Notice that every 4x4 switch is connected to another one by utilizing intra-quad lines which do not present intersections. A 4x4 Square Root accommodates four 2x2 quads (the blocks from A to D in Figure.8.4) and one central switch (the block indicated by E in the same Figure). Inter-quad express lanes are used for connecting each quad. Connections among them can be realized directly or through the central 2x2 quad. The inter-quad express lanes are affected by additional crossings which our modeling framework takes into account. Following the compositional approach, even a 8x8 Square Root topology can be obtained from four interconnected 4x4

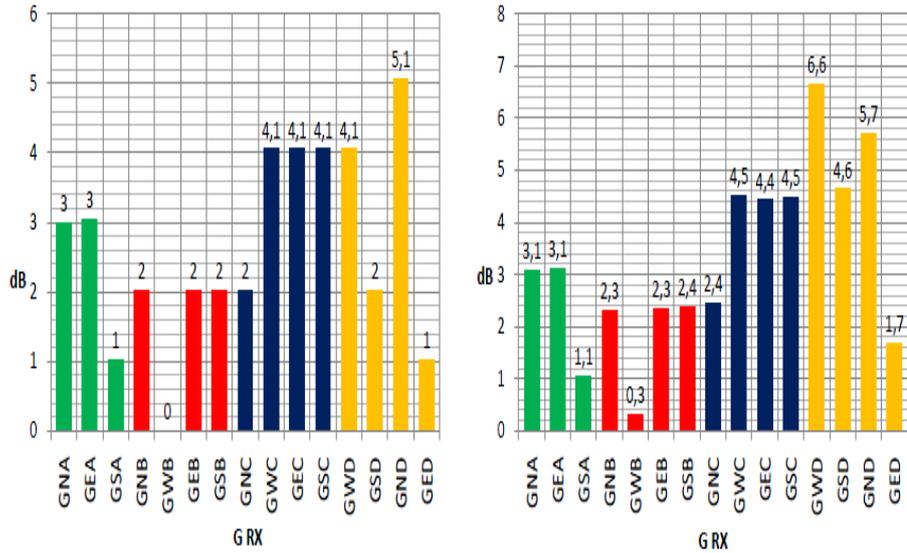


Figure 8.5: Insertion-Loss comparison for the 4x4 Square Root topology, by considering injection from G4 and with (right) and without (left) accounting for the inter-express lanes loss-contributions. Every intersection is optimized with standard elliptical tapers.

Square Roots. By using the previously obtained SystemC models of the different blocks, the insertion-loss analysis was accomplished. For brevity, only the case study with injection from Gateway G4 is reported.

Figure.8.5 (left) shows the insertion-loss from G4 to all other Gateways, without accounting for the inter-express lanes loss-contributions. The paths inside the Square Root topology are indicated by means of conventional acronyms. The abbreviation GNA stands for a communications between the Gateway G4 and the North-Eastern Gateway of quad A (G1), GED stands for the link between G4 and the South-Eastern Gateway of quad D (G14), and so on. In Figure.8.5 (right), on the contrary, these insertion losses have been taken into account, thus leading to a very accurate evaluation of the overall attenuation of each path and highlighting the importance of the losses for the overall performance of the network. Obviously, by changing from the first to

the second scenario, increasing of insertion-loss from 0.1 dB up to 2.6dB is observed. Our technology-annotated systemC models clearly enables to assess the physical properties of optical paths (i.e., a physical design issue), in addition to the traditional functional simulation capability. For the presented structure, the insertion-loss critical path (i.e. the path with the maximum amount of losses) is established from G4 to G14, and it reports 10dB when each crossing waveguide was designed by using the standard Elliptical Taper. By optimizing every crossing with the MMI-based solution, the insertion-loss is reduced to 4.85dB. Once determined the maximum insertion loss of the full network (which corresponds to the worst case optical path), it is possible to assess the minimum optical power that laser sources should provide to enable correct detection of the optical data stream at the photodetectors with the desired bit-error-rate (BER). By assuming for example that:

- (1) For a given bit-error-rate (BER) of  $10^{-9}$  the corresponding detector sensitivity is -20dBm;
- (2) Elliptically-tapered crossings are used at every intersection;
- (3) The data stream is carried over 30 different optical wavelengths;

The laser power injected in the network must be more than 3mW. This amount of power can be reduced by a factor of about 70% by optimizing each crossing with MMI-based structures, thus achieving 0.92mW. With the proposed approach based on the FDTD characterization of the fundamental blocks and a compositional approach in SystemC to form the complex network, the overall simulation times remain very low (fraction of seconds) and investigation of very complex network is certainly possible. In fact, long computation time is only required to characterize through FDTD the basic building blocks of the network, but this must be done only once for each block, when the corresponding s-parameters matrix should be determined, and does not need to be repeated each time the corresponding block is replicated inside the network.

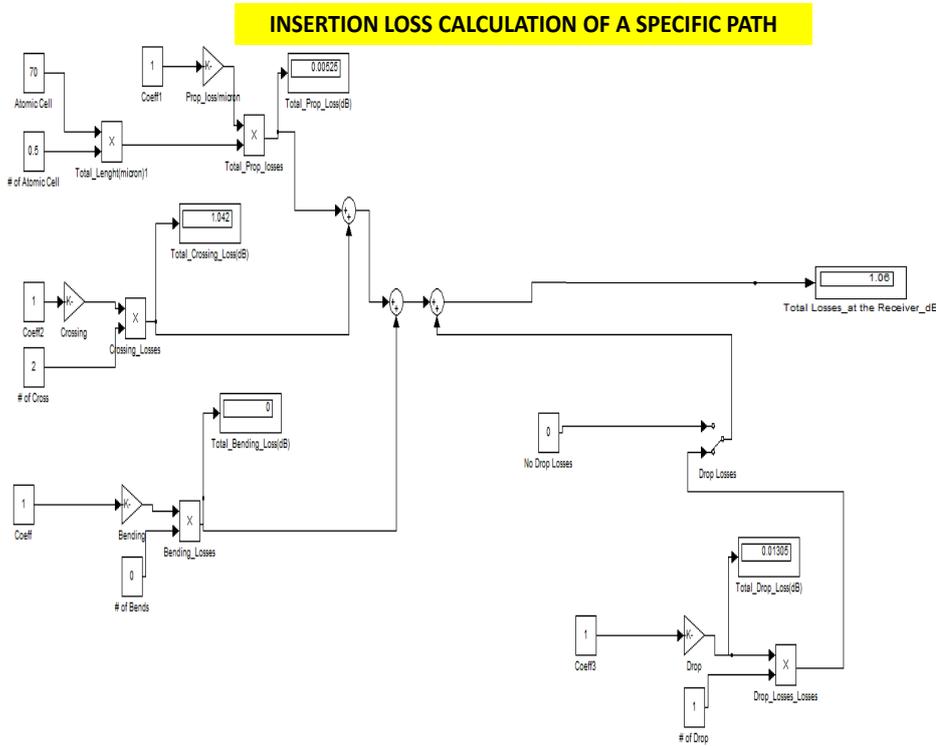


Figure 8.6: Insertion-Loss calculation of an optical path of a given optical NoC.

## 7 Simulink Simulation Framework for Optical NoCs

Simulink was also exploited to assess the insertion loss and latency results of any optical NoC examined in this thesis. In particular, all of physical parameters (loss and latency values) were imported into functional blocks such as straight, bend, crossing waveguides and also switching elements. By consequently compounding blocks among them, it was possible to model every path of the optical NoC under test, and finally assess all paths in terms of insertion loss and latency viewpoints. As illustrated in Figure.8.6, by summing up all of losses that the optical signal accrues along the communication path, the corresponding IL (Insertion Loss) of the ongoing path is obtained. In a similar way this methodology can be applied for latency evaluations. We opted for Simulink simulation framework because it offers high-level flexibility and high-speed simulation time, especially when high-complex ONoC

layouts comes into play. Although SystemC simulation framework is certainly an accurate network-level simulator, the higher complexity of physical layouts of ONoCs needs of a more flexible and practical tool able to cope with this limitation. Therefore, we leveraged on simulink simulation framework for the evaluation of the quality metrics of optical NoC physical layouts studied in this thesis.

## 8 Conclusion

This chapter has presented the modelling strategy based on the design-flow FDTD + SystemC, in order to explore and simulate optical network-on-chip topologies at system level. The procedure relies on the abstraction of the analytical models for the relevant components of an ONoCs (rings, waveguide crossings), toward the SystemC environment. The optical responses of the elementary photonic switching elements are first obtained via 2D-FDTD, then back-annotated in the SystemC modules. The modular and incremental composition of the basic switching elements allows the simulation of arbitrary complex topologies. A good trade-off between accuracy in the modeling of the investigated dynamics and simulation times has been demonstrated. Furthermore, the error parameters introduced to quantitatively validate the performance of the design-flow are largely satisfactory. As a case study, the SystemC modeling of one of the most famous optical network topologies, the 4x4-Square Root, is proposed and its insertion-loss assessment for different paths are quantified. Simulation results demonstrate that in the worst case insertion-loss is 10 dB using standard crossing optimization, while it is reduced at 4.85 dB using MMI-based intersections. Although SystemC simulation framework is certainly an accurate network-level simulator, the higher complexity of physical layouts of ONoCs needs of a more flexible and practical tool able to cope with this limitation. Therefore, Simulink simulation framework was adopted for the evaluation of quality metrics of optical NoC physical layouts studied in this thesis. Finally, this chapter has included contents that are referred to a cooperative and interdisciplinary work where further details are in [40].



# Thesis Conclusions

The presented Thesis aimed at searching for compelling cases that make silicon nanophotonic technology affordable in next generation many-cores systems. This was pursued by means of a trustworthy crossbenchmarking framework between an optical NoC and an aggressive electronic counterpart. The optimistic assessments of many previous works are put in discussion. Nonetheless, the experimental results do not paint a dismal picture on optical interconnect technology. In fact, it is proven to achieve relevant performance speedups even with bursty communication workloads (as opposed to the high utilization rates typically assumed), which are common in shared memory multiprocessors.

With conservative projections for optical component parameters, the major role played by static power is apparent. This calls for new power gating techniques. With more aggressive projections, the network interface turns out to be the clear bottleneck to achieve the break-even point with low-power ENOCs, hence it should be thoroughly analyzed for optimization.

When we extend the focus to the system as a whole, previous results paint a less dismal picture. By recalling that an interconnect fabric is in fact only a small portion of the total system energy, the ONoC is capable of speeding up the execution time meaning that the system as a whole can burn power for less time compared with the electrical counterpart. In light of this, an aggressive technology should not necessarily be adopted since energy savings are already there even with a conservative optical technology.



# Bibliography

- [1] Lee B.G et al “Ultrahigh Bandwidth silicon photonic nanowire waveguides for on chip networks”, IEEE Photon. Technol. Lett. 20 398-400.
- [2] Ophir N, et al, “Demonstration of 1.28 Tb/s transmission in next generation nanowires for photonics networks-on-chip”, 23rd Annual Meeting of the IEEE Photonic Society (Denver, CO,7-11, Nov, 2010) pp560-1.
- [3] Manipatruni S, Xu Q, Schmidt B, Shakya J, Lipson M, “High Speed carrier injection 18 Gbit/s silicon microring electro-optic modulator”,20th Annual meeting of the IEEE Lasers and Electro-Optics Society.pp537-8.
- [4] Preston K, Dong P, Schmidt B, Lipson M “High-speed all optical modulation using polycrystalline silicon microring resonators”, Appl. Phys. Lett. 92 151104.
- [5] Assefa S et al, “CMOS-integrated high-speed MSM germanium waveguide photodetector”. Opt.Express 184986-99.
- [6] Yin T et al, “Ge n-i-p waveguide photodetectors on silicon-on-insulator substrate”. Opt. Express. 15 13965-71.
- [7] Geis M W et al, “Silicon Waveguide infrared photodiodes with  $\lambda$  35 GHz bandwidth and phototransistors with 50 A/W-1 response”. Opt. Express 17 5193-204.
- [8] Jambois O, “Towards population inversion of electrically pumped Er ions sensitized by Si nanostructures” Opt. Express. 18 2230-5.
- [9] L. Eeckhout et al., “Designing Computer Architecture Research Workloads” Computer vol.36,pp.65-71.

- 
- [10] K. Hoste and L. Eeckhout, "Microarchitecture-Independent Workloads Characterization" *Micro,IEEE*,vol.27,pp.63-72,2007.
- [11] D. Wentzlaff et al , "On-Chip Interconnection Architecture of the Tile Processor", *IEEE Micro*, vol.27, n.5, pp. 15-31, November 2007.
- [12] H. Gu,K.H. Mo, and W. Zhang , "A low-power Low-cost Optical Router for Optical Network-on-Chip for Multiprocessor System-on-Chip ", *IEEE Computer Society Annual Symposium on VLSI*, 2009.
- [13] Delphine Marris-Morini et al., "HELIOS project: deliverable D010: state-of-the-art on Photonics on CMOS ", October 2011.
- [14] A. Biberman and Keren Bergman, "Optical Interconnection networks for high-performance computing systems", *Reports on Progress in Physics*, 2012.
- [15] B G. Lee, A Biberman, P Dong, M Lipson, K Bergman, " "All-Optical Comb Switch for Multiwavelength Message Routing in Silicon Photonic Networks", *IEEE Photonics Technology letters*,2008.
- [16] Ashok V. Krishnamoorthy "Photonics-to-electronics integration for optical interconnects in the early 21 st century", *Optoelectronics letters*,2006.
- [17] Pranay Koka et al. "Silicon-Photonic Network Architectures for Scalable, Power-Efficient Multi-Chip Systems", *ISCA 2010: International Symposium on Computer Architecture*, June 2010.
- [18] David A.B. Miller, "Rationale and Challenges for Optical Interconnects to Electronic Chips", 2000.
- [19] ePIXfab "The silicon photonics platform".
- [20] I. Yoiung, E. Mohammed, J. Liao, A. Kern, S. Palermo, B. Block, M. Reshotko, P. Chang, "Optical I/O technology for tera-scale computing". *ISSCC, Dig.Tech.Papers*, pp. 468-469, Feb 2009.
- [21] F. Liu et al, "10 Gbps 530fJ/b Optical Transceiver Circuit in 40nm CMOS" *Symposium on VLSI Circuit*,pp.290-291, June 2011.

- 
- [22] A. Huang et al, “A 10Gb/s photonic modulator and WDM MUX/DEMUX integrated with electronics in 0.13um SOI CMOS” ISSCC Dig.Tech.Papers,pp922-929,Feb.2006.
- [23] J. Orcutt et al, “Open foundry platform for high-performance electronic-photonic integration” *Optics Express*,vol.20, no 11, pp 12222-12232, May 2012.
- [24] A. Biberman et al., “Photonic Network-on-Chip Architectures Using Multilayer Deposited Silicon Materials for High-Performance Chip Multiprocessors”, *ACM JETC*, Vol.7, Issue 2, pp. 7:1-7:25, 2011
- [25] H. Wang et al., “Nanophotonic Optical Interconnection Network Architecture for On-Chip and Off-Chip Communications”, *OFC/NFOEC’08: Conference on Optical Fiber communication/National Fiber Optic Engineers*, May 2008.
- [26] 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.
- [27] 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.
- [28] X. Tan et al., “On a Scalable, Non-Blocking Optical Router for Photonic Networks-on-Chip Designs”, *Photonics and Optoelectronics (SOPO)*, May 2011.
- [29] M. J. Cianchetti, J. C. Kerekes, and D. H. Albonesi “Phastlane: A rapid transit optical routing network”, *ISCA’09: International Symposium on Computer Architecture*, June 2009.
- [30] S. Beamer et al., “Re-Architecting DRAM Memory Systems with Monolithically Integrated Silicon Photonics”, *ISCA’10: International Symposium on Computer Architecture*, June 2010.

- 
- [31] 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.
- [32] D. Vantrease et al., “Corona: System implications of emerging nanophotonic technology”, *ISCA’08: International Symposium on Computer Architecture*, June 2008.
- [33] Y. Pan, P. Kumar, J. Kim, G. Memik, Y. Zhang, and A. Choudhary *Firefly: Illuminating future network-on-chip with nanophotonics?*, *Intl. Symp. on Computer Architecture*, pages 429440, Austin, TX, June 2009.
- [34] C. Batten et al. *Designing chip level nanophotonic interconnection networks* *Emerging and Selected Topics in Circuits and Systems*, *IEEE Journal*. vol. 2, no. 2, pp. 137-153, 2012.
- [35] 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.
- [36] L. Ramini, P. Grani, Herve F. Tatenguem, A. Ghiribaldi, S. Bartolini and D. Bertozzi “Assessing the Energy Break-Even Point between an Optical NoC Architecture and an Aggressive Electronic Baseline”, *Proceedings of the Conference on Design, Automation and Test in Europe*, (DATE 14), Dresden,(Germany), March, 2014.
- [37] L. Ramini and D. Bertozzi “Power Efficiency of Wavelength-Routed Optical NoC Topologies for Global Connectivity of 3D Multi-core Processors”, *NoCarch’12: International Workshop on Networks-on-Chip Architecture*, December 2012.
- [38] L. Ramini, P.Grani, S. Bartolini and D. Bertozzi “Contrasting Wavelength-Routed Optical NoC Topologies for Power Efficient 3D-Stacked Multicore Processors Using Physical Layer Analysis”, *DATE’13: International Conference on Design Automation and Test in Europe*, March 2013.

- [39] A. Boos, L. Ramini, U.Schlichtmann and D. Bertozzi “PROTON: An Automatic Place-and-Route Tool for Optical Network-on-Chip”, IC-CAD’13: International Conference on Computer Aided Design, November 2013.
- [40] A. Parini, L. Ramini, F. Lanzoni ,G. Bellanca and D. Bertozzi “Bottom-Up Abstract Modelling of Optical Networks-on-Chip: From the Physical to the Architectural Layer”, International Journal of Optics, November, 2012.
- [41] I. O’Connor et al., “Towards Reconfigurable Optical Networks on Chip”, ReCoSoC 2005, pp.121-128.
- [42] David Wentzlaff,Patrick Griffin,Henry Hoffmann,Liewei Bao,Bruce Edwards,Carl Ramey,Matthew Mattina,Chyi-Chang Miao,John F. Brown III and Anant Agarwal “On-Chip Interconnection Architecture of the Tile Processor”, IEEE computer society online at [www.computer.org/join/](http://www.computer.org/join/).
- [43] J. U. Knickerbocker et al, “Three-dimensional silicon integration”, IBM J.RES.& DEV., vol. 52 NO. 6, November, 2008.
- [44] International Technology Roadmap for Semiconductors (ITRS). <http://www.itrs.net/>.
- [45] D. Bertozzi, A. Jalabert, S. Murali, R.Tamhankar, S. Stergiou, L. Benini, and G. De Micheli, “NoC synthesis ow for customized domain specific multiprocessor systems-on-chip”, IEEE Transactions on Parallel and Distributed Systems, vol. 16, n. 2, February 2005, pp.113-128.
- [46] J. Chan, G. Hendry, A. Biberman, K.Bergman, L.P. Carloni, ”PhoenixSim: A Simulator for Physical-Layer Analysis of Chip-Scale Photonic Interconnection Networks”, Design Automation and Test in Europe Conference 2010, pp.691-696.
- [47] E. Drouard, M. Briere, F. Mieyeville, I.O’Connor, X. Letartre, F. Gaot, “Optical Network on chip Multi-Domain modelling using SystemC”. FDL 2004, pp.123-135.

- 
- [48] M. Briere et al., "Heterogeneous Modelling of an Optical Network-on-Chip with SystemC", IEEE Int. Workshop on Rapid System Prototyping, pp.10-16, 2005.
- [49] S. Yee, "Numerical Solution of Initial Boundary Value Problems Involving Maxwell's Equations in Isotropic Media", IEEE Trans. on Antennas and Propagat., vol. 14, May 1966, pp.302-307.
- [50] Otto Schwelb, "Transmission, Group Delay, and Dispersion in Single-Ring Optical Resonators and Add/Drop Filters-A Tutorial Overview", Journal of Lightwave Technology, vol. 22, n. 5, May 2004.
- [51] Jun So Pak, Chunghyun Ryu, Joungho Kim, "Multi-Stacked Through-Silicon-Via Effects on Signal Integrity and Power Integrity for Application of 3-Dimensional Stacked-Chip-Package", IEEE Micro, 1998, 18, (4), pp.17-22.
- [52] A. Scandurra and I.O'Connor, "Scalable CMOS-compatible photonic routing topologies for versatile networks on chip", Network on Chip Architecture, 2008.
- [53] 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.
- [54] 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.
- [55] 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.
- [56] 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.

- [57] S. Le Beux et al., "Reduction Methods for adopting network on chip topologies to 3D Architectures", Book on Microprocessor and Microsystem, volume 37, issue 1, 2012.
- [58] Nevin Kirman Jos'e F. Mart'inez "A Power-efficient All-optical On-chip Interconnect Using Wavelength-based Oblivious Routing", ISCA, 2010.
- [59] H. Gu et al., "A low-Power fat tree-based Optical Network-on-Chip for Multiprocessor System-on-Chip", DATE'09: Conference on Design, Automation and Test in Europe, 2009.
- [60] 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.
- [61] F. Xu and A.W Poon, "Silicon cross-connect filters using microring resonator coupled multimode interference-based waveguide crossings", Optics Express, 2008.
- [62] Xuezhe Zheng et al., "Ultra-efficient 10Gbit/s hybrid integrated silicon photonic transmitter and receiver", Opt Express, 14;19(6):5172-86, March 2011.
- [63] David A.B. Miller, "Energy consumption in optical modulators for interconnects", Opt Express, Vol. 20, pp. A293-A308, March 2012.
- [64] G.P. Agrawal, "Fiber-Optic Communication Systems", Wiley-Interscience, third edition, chapter fourth, pp.133-178, 2002.
- [65] Somayyeh Koochi, Meisam Abdollahi, Shaahin Hessabi "All-Optical Wavelength-Routed NoC based on a Novel Hierarchical Topology", NOCS, 2011 Pittsburgh, PA, USA.
- [66] S.Koochi, S.Hessabi, S.J.B.Yoo "An Optical Wavelength Switching Architecture for a High-Performance Low-Power Photonic Network on Chip", Advanced Information Networking and Applications (WAINA), 2011 IEEE Workshops of International Conference.

- [67] M.Georgas et al., “A Monolithically-Integrated Optical Receiver in Standard 45-nm SOI”, Solid State Circuits,2002.
- [68] S. Stergiou, F. Angiolini, S. Carta, L. Raffo, D. Bertozzi, G. De Micheli. *timespipes lite: a synthesis oriented design library for networks on chips*. Design, Automation and Test in Europe, 2005. pp. 1188-1193 Vol. 2.
- [69] A. Udipi, N. Muralimanohar, R. Balasubramonian, A. Davis, N. Jouppi. *Combining memory and a controller with photonics through 3d-stacking to enable scalable and energy efficient systems*. ISCA 11.
- [70] A. Hansson et al. *Avoiding messagedependent deadlock in network-based systems on chip*. VLSI Design, vol. 2007. 2007
- [71] F. Gilabert, M.E. Gomez, S. Medardoni, D. Bertozzi. *Improved utilization of noc channel bandwidth by switch replication for cost-effective multi-processor systems-on-chip*. (NOCS), 2010. ACM/IEEE International Symposium on, 2010, pp. 165-172.
- [72] C. Batten, A. Joshi, V. Stojanovic, K. Asanovic. *Designing chiplevel nanophotonic interconnection networks*. Emerging and Selected Topics in Circuits and Systems, IEEE Journal. vol. 2, no. 2, pp. 137-153, 2012
- [73] S.Murali et al, “Sunfloor:Application-Specific Desing of Neetworks-on-Chip.”, Desing Automation and Test in Europe,2006
- [74] A.Biberman et al, “High-Speed Data Trasmission in Multi-Layer Deposited Silicon Photonics for Advanced Photonic Networks-on-Chip., CLEO 2011CThA1, 2011
- [75] IPOPT package,<https://projects.coin-or.org/Ipopt>,2012
- [76] N. Binkert, R.G. Dreslinski, L.R. Hsu, K.T. Lim, A.G. Saidi, S.K. Reinhardt *The m5 simulator: Modeling networked systems*. Micro, IEEE, vol. 26, no. 4, pp. 52-60, 2006.
- [77] C. Bienia, S. Kumar, J.P. Singh, K. Li. *The parsec benchmark suite: characterization and architectural implications*. PACT 08. <http://doi.acm.org/10.1145/1454115.1454128>.

- [78] H. Falaki, R. Mahajan, S. Kandula, D. Lymberopoulos, R. Govindan, D. Estrin. *Diversity in smartphone usage*. Proceedings of the 8th international conference on Mobile systems, applications, and services. <http://doi.acm.org/10.1145/1814433.1814453>
- [79] A. Strano, D. Ludovici, and D. Bertozzi. “A library of dual-clock fifos for cost-effective and flexible mp soc design”. In Int. Conf. on Embedded Computer Systems (SAMOS), pages 20–27, 201.
- [80] W. Dally and B. Towles. “Principles and Practices of Interconnection Networks.” Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 2003.
- [81] J. Leu, V. Stojanovic. *Injection-locked clock receiver for monolithic optical link in 45nm SOI*. (A-SSCC), 2011 IEEE Asian, 2011, pp. 149-152.
- [82] M.Ortín, L.Ramini, Herve F. Tatenguem, V.Vgnals and Davide Bertozzi. *A Complete Electronic Network Interface Architecture for Global Contention-Free Communication over Emerging Optical Networks-on-Chip*. (A-SSCC), Proceedings of the International Symposium (GLSVLSI 14), Huston,(USA), May, 2014.
- [83] Naveen Moralimanohar and Rajeev Balasubramonian, “CACTI 6.0: A toll to model large caches”, IEEE MICRO,2006.
- [84] G.R.Hadley, “Effective index model for vertical-cavity surface-emitting lasers”, Opt.Lett.,vol.20,pp.1483-1485,1995.



# Publications

## Conference:

- (1) A. Parini, **L. Ramini**, G. Bellanca, D. Bertozzi: *Abstract Modelling Switching Elements for Optical Network-on-Chip with Technology Platform Awareness*. Proceedings of the Fifth ACM Interconnection Network Architecture, On-Chip Multi-Chip Workshop (INA-OCMC 11), pp. 31-33, Heraklion, (Greece), January, 2011.
- (2) **L. Ramini**, D. Bertozzi, L P. Carloni: *Engineering a Bandwidth Scalable Optical Layer for a 3D Multi-core Processor with Awareness of Layout Constraints*. Proceedings of the Sixth IEEE Symposium on Network-on-Chip, (NOCS 12), Lyngby, (Denmark), May, 2012.
- (3) **L. Ramini** and D. Bertozzi: *The Design Predictability Concern in Optical Networks-on-Chip Design*. Proceedings of the ACP Conference, Asia Communication and Photonics Conference. ACP 12), Canton, (China), November, 2012.
- (4) **L. Ramini** and D. Bertozzi: *Power Efficiency of Wavelength-Routed Optical NoC Topologies for Global Connectivity of 3D Multi-Core Processors*. Proceedings of the Networks-on-Chip Architectures Workshop, (NoCArC 12). Vancouver, (Canada), December, 2012.
- (5) **L. Ramini**, P. Grani, S. Bartolini and D. Bertozzi: *Contrasting Wavelength-Routed Optical NoC Topologies for Power-Efficient 3D-stacked Multicore Processors using Physical-Layer Analysis*. Proceedings of the Design, Automation and Test in Europe, (DATEâ13), Grenoble, (France), March, 2013.
- (6) A. Boos, **L. Ramini**, U. Shiltzmann and D. Bertozzi: *PROTON: An Automatic Place-and-Route Tool for Optical Network-on-Chip*. Proceedings of the International Conference on Computer Aided Design , (ICCAD 13),

San Jose, (USA), November 2013.

(7) **L. Ramini** and D. Bertozzi : *Crossbenchmarking an optical network-on-chip with an aggressive electrical baseline with physical layer awareness.*

Proceedings of 3rd International Symposium on Photonics and Electronics Convergence, (ISPEC 13), Tokyo, (Japan), November, 2013.

(8) **L. Ramini**, P. Grani, Herve F. Tatenguem, A. Ghiribaldi, S. Bartolini and D. Bertozzi: *Assessing the Energy Break-Even Point between an Optical NoC Architecture and an Aggressive Electronic Baseline*. Proceedings of the Conference on Design, Automation and Test in Europe, (DATE 14), Dresden,(Germany), March, 2014.

(9) M.Ortín, **L.Ramini**, V.Vgnals and Davide Bertozzi: *Capturing the Sensitivity of Optical Network Quality Metrics to its Network Interface Parameters.* Proceedings of the HiPEAC 1st International Workshop on Exploiting Silicon Photonics for energy-efficient heterogeneous parallel architectures (SiPhotonics'2014), Vienna, (Austria), January, 2014.

(10) M.Ortín, **L.Ramini**, Herve F. Tatenguem, V.Vgnals and Davide Bertozzi: *A Complete Electronic Network Interface Architecture for Global Contention-Free Communication over Emerging Optical Networks-on-Chip.* Proceedings of the International Symposium (GLSVLSI 14), Huston,(USA), May, 2014.

## Journal & Book Chapters:

(11) A. Parini, **L. Ramini**, F. Lanzoni ,G. Bellanca and D. Bertozzi: *Bottom-Up Abstract Modelling of Optical Networks-on-Chip: From the Physical to the Architectural Layer*. International Journal of Optics, November, 2012.

(12) [in press], **L. Ramini**, and D. Bertozzi, CRC Emerging VLSI Systems, 2014.: *Design Space Exploration of Wavelength-Routed Optical NoC Topologies for a 3D-Stacked Multi-nad-Many Core Processors.*

## Invited Talk:

Luca Ramini was invited speaker and track chairman at the System-on-Chip Conference, Irvine, (USA), October 2013. His presentation, *Optical Interconnect technology for 3D Stacked Multi- and-Many-core Systems: When, Where and How* was included in the SoC conference's proceeding.

# Acknowledgements

First of all, I would like to thank my future wife Lisa and all my family. I would like to thank my supervisor, the Professor Davide Bertozzi who always believed in me during this experience. Also, I thank very much all of partners of the Photonica project, the Prof. Gaetano Bellanca (University of Ferrara, Italy), the Prof. Giovanna Calò (Politecnico of Bari, Italy), as well as the Prof. Sandro Bartolini (University of Siena, Italy).

One special thanks goes to the Professor Luca Carloni (Columbia University, NYC, USA) who kindly hosted me for six months at Columbia University, during the first year of my Ph.D. I will never forget all of friends and colleagues that I had the pleasure to work with them over these unforgettable years. Thank you nonno for giving me the strengths of living this experience with love and passion as you made in your life.