Multi-Optical Network-on-Chip for Large Scale MPSoC

Sébastien Le Beux, Jelena Trajkovic, Ian O'Connor, Gabriela Nicolescu, Guy Bois, Pierre Paulin

► To cite this version:

Sébastien Le Beux, Jelena Trajkovic, Ian O'Connor, Gabriela Nicolescu, Guy Bois, et al.. Multi-Optical Network-on-Chip for Large Scale MPSoC. IEEE Embedded Systems Letters, 2010. inria-00618593

HAL Id: inria-00618593

<https://hal.inria.fr/inria-00618593>

Submitted on 2 Sep 2011

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.

Multi-Optical Network-on-Chip for Large Scale MPSoC

Sebastien Le Beux, Jelena Trajkovic, Ian O'Connor, Gabriela Nicolescu, Guy Bois, and Pierre Paulin

Abstract—Optical network-on-chip (ONoC) architectures are emerging as promising contenders to solve bandwidth and latency issues in multiprocessor systems-on-chip (MPSoC). However, current on-chip integration technologies for optical interconnect allow interconnecting only dozens of IPs. Scaling with MPSoCs composed of hundreds of IPs thus, relies on unpredictable technological innovations. In this letter, we propose a method that combines multiple ONoCs. Each ONoC is small enough to rely on already existing and proven technologies. We evaluate the approach for various interconnect scenarios, showing that it scales well with the size of the MPSoC architectures.

Index Terms—Design methodology, optical network-on-chip (ONoC).

I. INTRODUCTION

THE shift to very high performance distributed multiprocessor systems-on-chip (MPSoC) is the recognized route to reach the performance requirement for compute-intensive applications [2]. MPSoC requires high-speed communication between processors, which clearly relies upon the existence of a fast and flexible interconnect network. Optical networks-on-chip (ONoC) overcome the limitations of electrical interconnect through the use of wavelength division multiplexing (WDM) [9], which allows concurrent transition of multiple optical signals on the same waveguide. However, the reliability of such optical communication is driven by the laser output power and the coupling losses in optical switches. Thus, the scalability of the ONoC relies on technological advances in optical on-chip interconnect.

Most of the previous work contrary to ours, uses off-chip lasers sources [11], [3]–[5]. We consider on-chip laser sources due to two reasons: on-chip lasers can be better integrated and consume less power [14], leading to smaller system dimension and lower power consumption. However, output power of on-chip lasers is lower than that of the off-chip ones, which may

imply more stringent constraints on propagation losses. Nevertheless, on-chip laser sources are more suitable for low-power systems. Due to the constantly increasing number of IPs in MPSoC, the power budget of the systems with off-chip lasers may also become constrained. Regardless of the topology, the number of IPs connected through an ONoC directly impacts the number of the required optical switches and the losses leading to unfeasible designs. Therefore, we propose to connect large scale MPSoC through multi-ONoC. This method divides a large, infeasible ONoC into multiple, smaller ONoCs. In the resulting multi-ONoC architecture, each ONoC is reliable and feasible. As drawbacks, additional electrical routing is required and the number of waveguides increases, leading to a more complex layout on the optical layer. This method is evaluated on a multistage network architecture, but it can be adapted to others topologies.

Several papers have addressed ONoC design issues exploiting both electrical and optical NoC technologies. In [4], electrical interconnect is used for control flow and the optical is used for the data flow. The electrical signal precedes the optical one and reserve an optical path. Therefore, optical communications may be delayed until an optical path becomes free, resulting in contention delay. The same applies to the fat tree [5] and the mesh [6] ONoCs. Contrary to the above mentioned, the ONoC we are considering is contention free [7], [1]. In [3], electrical interconnects manage local communication, while optical ones are responsible for global communication. In the chosen ring topology, each wavelength flowing through the ONoC must be assigned to a given optical network interface, disallowing parallel communications through the same wavelength, where our architecture allows parallel communications using the same wavelength. Complementary to our work, [12] and [13] consider layout issues: [12] introduces a simulator for optical on-chip interconnects, and [13] a tool for placement and routing. Reference [11] also addresses scalability issues, but uses off-chip laser sources.

II. SINGLE-ONoC ARCHITECTURE

An ONoC interconnects a set of computing and storage resources that are both considered as IPs. As detailed in [7], ONoC interconnects the IPs according to a binary IP connectivity matrix: a 1 means that a communication is possible from a source IP to a target IP. In the connectivity matrix illustrated in Fig. 1(b), eight IPs are considered and each IP can communicate with all other IPs, except with itself. Each IP is connected to the ONoC through optical network interface (ONI) which perform electro-optical and opto-electrical conversions. ONI can communicate simultaneously with one or several other ONIs through passive

Manuscript received February 05, 2010; revised April 12, 2010; accepted June 21, 2010. Date of publication July 15, 2010; date of current version September 17, 2010. This manuscript was recommended for publication by D. Atienza.

S. Le Beux, J. Trajkovic, G. Nicolescu, and G. Bois are with the Département de Génie Informatique et de Génie Logiciel, École Polytechnique de Montréal, Montréal, QC H3C 3A7, Canada (e-mail: sebastien.le-beux@polymtl.ca; jelena.trajkovic@polymtl.ca; gabriela.nicolescu@polymtl.ca; guy.bois@polymtl.ca).

I. O'Connor is with the Lyon Institute of Nanotechnology (INL), École Centrale de Lyon, Ecully, 69130, France (e-mail: ian.oconnor@ec-lyon.fr).

P. Paulin is with the STMicroelectronics, Ottawa, ON K2H 8R6, Canada (e-mail: pierre.paulin@st.com).

Color versions of one or more of the figures in this paper are available online at <http://ieeexplore.ieee.org>.

Digital Object Identifier 10.1109/LES.2010.2057407

Fig. 1. Baseline ONoC a) wavelength matrix, b) IP connectivity matrix, and c) topology schematic.

photonic routing structure: λ -router. A communication through the ONoC is performed in three stages: 1) the source ONI converts an electrical to an optical signal; 2) the optical signal is routed through the λ -router; and 3) the target ONI converts the optical signal into an electrical one.

The baseline ONoC architecture [7] is a λ -router network that consists of N stages of *optical switches*. Each stage is composed of optical switches characterized by the same resonant wavelength and linked through *waveguides*. The waveguides transmit optical signals and the optical switches manage routing by selecting and redirecting an input signal based on its wavelength. The path followed by the optical signal in the λ -router depends only on the wavelength. WDM technique is used: when multiple signals of various wavelengths are injected at the input, where individual signals simultaneously obey the routing characteristics of the optical switch according to their individual wavelengths. Because of this property, the presented ONoC is contention-free.

Only one physical path associated with a single wavelength exists between a source ONI and a target ONI. A wavelength to be used for each communication is specified in a wavelength matrix. This matrix is for a full connectivity between all ONIs, and for a desired connectivity it needs to be superposed to a desired connectivity matrix. Fig. 1(a) represents the initial wavelength matrix for an ONoC of eight ONIs. The values that are used after superposing with the desired connectivity matrix from the Fig. 1(b) are emphasized. From these connections, the ONoC is built following the reduction method proposed in [7]. This method adapts ONoC for a given IP connectivity and reduces the implementation complexity of the optical interconnect by suppressing unused optical switches. The resulting ONoC is illustrated in Fig. 1(c).

The ONoC design feasibility and efficiency is constrained by the propagation losses occurring in the passive photonic components. For the current technology, optical switches and waveguides, respectively, introduce 0.3 dB and 2 dB/cm losses. To achieve acceptable communications bit error rate, the number of optical switches crossed (nOSC) and the length of waveguides must be limited accordingly to the output power of the laser source. For instance, a proposed integrated laser output power is $2.5 \mu\text{W}/\mu\text{m}^2$ [8], which allows crossing of 48 optical switches and 1 cm waveguide. However, interconnecting hundreds of IPs through a single-ONoC relies on unpredictable advances in optical on-chip interconnect integration technologies. Therefore, we propose a methodology that scales with large optical interconnect requirements according to proven technology.

Fig. 2. Multi-ONoC architecture interconnecting eight IPs through n ONoCs.

III. MULTI-ONoC ARCHITECTURE

In order to interconnect large MPSoC, multiple (feasible) ONoCs can be combined. The considered architecture, illustrated in Fig. 2, is composed of a set of IPs interconnected to multi-ONoC through electrical network interfaces (ENIs). Each ONoC manages part of the connections specified in the IP connectivity matrix. As a consequence, the packet routing is realized in three steps: 1) electrical routing from a source IP to the relevant ONoC (based on the target IP address); 2) optical routing through the ONoC; and 3) electrical routing from the ONoC to the target IP. Electrical routing is realized by address decoder components.

The connections specified in the IP connectivity matrix are distributed so that ONoCs manage an equivalent number of wavelengths and each stage of the λ -router is associated with a single wavelength. By supporting only a subset of wavelengths, only a subset of stages is required, which reduces the nOSC. This is implemented via the reduction method described in Section II.

To illustrate this approach, we start from the example shown in Fig. 1, where a single-ONoC is used to interconnect eight IPs. We transform a single-ONoC into a multi-ONoC, where ONoC₁ manages the connections using λ_1 – λ_4 , and ONoC₂ manages the connections using λ_5 – λ_8 . The resulting connectivity matrices and multi-ONoC architecture are shown in Fig. 3. Compared to the starting scenario (Fig. 1), the nOSC decreases

Fig. 3. a) ONoC₁, b) ONoC₂ connectivity matrixes, and c) the resulting multi-ONoC architecture.

Fig. 4. Modified ONI location: a) scemaic and b) layout.

from seven to four, which satisfies design constraints. As drawbacks, the number of required waveguides increases and address decoders are necessary to multiplex and demultiplex the data. This may lead to additional constraints in the layout and additional latencies in communication. Also, the overall cost of ONIs slightly increases due to the addition of address decoders. Note that the total number of the laser sources stays the same as in the initial, single-ONoC scenario.

The layout shown in Fig. 3(c) implies that waveguides necessarily cross each other, resulting in additional coupling losses (0.05 dB per waveguide crossing [10]). This drawback limits the benefits of reducing the number of optical switches crossed. However, since each ONoC may occupy a part of a die (2-D architecture) or an optical layer (3-D architecture), a simple modification of the ONI location in the layout allows removing these extra waveguide crossing (Fig. 4). In this example, the locations of ONI receivers and ONI emitters of the ONoC₁ (from Fig. 2) are modified. Since the layout suggested in Fig. 4(a) exhibits regularity, there is potential for automatic layout generation. A minor drawback to ONI location modification may be a slight increase in the die size to allow for placement and routing of the additional waveguides.

TABLE I

COMPARISON BETWEEN DESIGN METRICS FOR VARIOUS NUMBER OF ONoCs

Number of ONoCs	1	2	3	4	8
Number of waveguides	8	16	24	32	56
nOSC	8	4	3	2	1

Table I shows nOSC and number of waveguides for one, two, four, and eight ONoCs scenarios interconnecting eight IPs. When the number of ONoCs is increased, the nOSC decreases and the number of waveguides increases.

IV. METHOD AND EXPERIMENTAL RESULTS

A. Method

The multi-ONoC generation method is given in Algorithm 1. As an initial configuration, the number of ONoCs is set to one since it maximizes WDM benefits. The inputs of the method are a connectivity matrix for N IPs and the maximum tolerated nOSC. In step 3), the set of used wavelengths (WL) is equally distributed to the ONoCs. From this distribution, in step 4), the connectivity of each ONoC is obtained and the reduction method suppresses the unused optical switches. The maximum nOSC of the resulting ONoCs is compared to the tolerated nOSC. In the case where the constraint is satisfied, the multi-ONoC configuration is validated and the process terminates. Otherwise, the number of ONoCs increases and the process is repeated. By adapting steps 3) and 4) of the algorithm, our methodology can also be applied to different ONoC topologies.

Algorithm 1: Multi-ONoC Generation

- 1) Set the number of ONoC to 1.
- 2) Get the tolerated nOSC and the connectivity matrix.
- 3) Equally distribute wavelengths $wl_i \in WL$ to the ONoCs.
- 4) For each ONoC, generate the connectivity matrix and call the reduction method.
- 5) Set nOSC as the maximum nOSC of all the resulting ONoCs.
- 6) If nOSC satisfies the design constraint, validate the multi-ONoC and exit; else, increment the number of ONoC and go to step 3).

B. Experimental Results

Fig. 5 shows how the proposed method enables feasible interconnect of large scale MPSoC using state-of-the-art on-chip optical interconnect integration technology. We use the full connectivity to connect 128, 256, and 512 IPs, and set the maximum tolerated nOSC to 48. The number of required ONoCs increases with the number of interconnected IPs: 128, 256, and 512 IPs require three, six, and 11 ONoCs, respectively. Also, the number of required waveguides increases linearly with the number of required ONoCs.

Next, we evaluate the efficiency of our method for various connectivity matrices. They represent realistic scenarios where sets of IPs are clustered so that they do not need to communicate

Fig. 5. Scalability of the method: a) nOSC and b) number of waveguides.

Fig. 6. Efficiency of the method for interconnecting 288 clustered IPs: a) nOSC and b) number of waveguides.

with each other through the ONoC, like when a cluster of processors is connected to a cluster of memories. The number of connections in the connectivity matrix will vary depending on the number of clusters. We equally distributed 288 IPs into two, three, four, and six clusters. According to Fig. 6(a), six ONoCs

are necessary to satisfy the design constraint. Our method effectively reduces the nOSC for all of the considered number of clusters. The only exception is the two-cluster architecture, which is already strongly optimized for one ONoC. The number of waveguides required to implement the multi-ONoC depends on the number of clusters [Fig. 6(b)]. While six ONoCs are required for all the considered architecture configurations, the complexity (the number of waveguides) is reduced with the number clusters.

V. CONCLUSION

Interconnecting a large number of IPs using optical networks is infeasible due to the current technological constraints. We propose a methodology that enables ONoC designers to overcome this challenge by dividing a large ONoC into multi-ONoCs. At the acceptable price of extra waveguides and electrical routing, the proposed method reduces the number of nOSC, making it possible to interconnect large number of IPs. In our future work, we will evaluate the benefits of this methodology to different topologies, such as ring.

REFERENCES

- [1] A. Kazmierczak *et al.*, "Design, simulation, and characterization of a passive optical add-drop filter in silicon-on-insulator technology," *IEEE Photon. Technol. Lett.*, vol. 17, pp. 1447–1449, Jul. 2005.
- [2] International Technology Roadmap for Semiconductors ITRS [Online]. Available: <http://public.itrs.net/>
- [3] S. Pasricha *et al.*, "ORB: An on-chip optical ring bus communication architecture for multi-processor systems-on-chip," in *Proc. Asia South Pacific Des. Autom. Conf.*, Seoul, Korea, 2008, pp. 789–794.
- [4] A. Shacham *et al.*, "Photonic networks-on-chip for future generations of chip multiprocessors," *IEEE Trans. Comput.*, vol. 57, no. 9, pp. 1246–1260, Sep. 2008.
- [5] H. Gu *et al.*, "A low-power fat tree-based optical network-on-chip for multiprocessor system-on-chip," in *Proc. Des., Autom., Test Eur.*, Nice, France, 2009, p. 3.
- [6] H. Gu *et al.*, "A novel optical mesh network-on-chip for gigascale systems-on-chip," in *Proc. Asia South Pacific Conf. Circuits Syst.*, Macao, China, 2008, pp. 1728–1731.
- [7] I. O. Connor *et al.*, "Reduction methods for adapting optical network on chip topologies to specific routing applications," in *Proc. Des. Circuits Integrated Syst.*, Grenoble, France, 2008.
- [8] T. Spuesens *et al.*, "Improved design of an InP-based microdisk laser heterogeneously integrated with SOI," in *Proc. 6th IEEE Int. Conf. Group IV Photon.*, San Francisco, CA, Sep. 2009, pp. 202–204.
- [9] G. Rouskas, *Routing and Wavelength Assignment in Optical WDM Networks*. New York: Wiley, 2001, Encyclopedia of Telecommunications, DOI: 10.1002/0471219282.EOT205.
- [10] T. Fukazawa *et al.*, "Low loss intersection of Si photonic wire waveguides," *Jap. J. Appl. Phys.*, vol. 43, no. 3, pp. 646–647, 2004.
- [11] Y. Pan *et al.*, "Firefly: Illuminating future network-on-chip with nanophotonics," *SIGARCH Comput. Archit. News*, vol. 37, p. 3, 2009.
- [12] J. Chan *et al.*, "PhoenixSim: A simulator for physical-layer analysis of chip-scale photonic interconnection networks," in *Proc. Des., Autom., Test Eur.*, Grenoble, France, 2010, pp. 691–696.
- [13] D. Ding *et al.*, "O-Router: An optical routing framework for low power on-chip silicon nano-photonic integration," in *Proc. Des. Autom. Conf.*, San Francisco, CA, 2009, pp. 264–269.
- [14] J.-M. Fedeli *et al.*, "Integration issues of a photonic layer on top of a CMOS circuit," in *Proc. SPIE Photon. Eur.*, Apr. 3–7, 2006, pp. 97–111.