Optical Ring Network-on-Chip (ORNoC): Architecture and Design Methodology

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.. Optical Ring Network-on-Chip (ORNoC): Architecture and Design Methodology. Design, Automation & Test in Europe Conference & Exhibition (DATE), Mar 2011, Grenoble, France. inria-00618600

HAL Id: inria-00618600

<https://hal.inria.fr/inria-00618600>

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.

Optical Ring Network-on-Chip (ORNoC): Architecture and Design Methodology

Sébastien Le Beux^{*§}, Jelena Trajkovic[§], Ian O'Connor^{*}, Gabriela Nicolescu[§], Guy Bois[§] and Pierre Paulin[‡]

[§]Software and Computer Department
Ecole Polytechnique de Montréal, Canada
Email: {Jelena.Trajkovic, Gabriela.Nicolescu, Guy.Bois}@polymtl.ca

^{*}Lyon Institute of Nanotechnology,
Ecole Centrale de Lyon, France
Email: {Sebastien.Le-Beux, Ian.OConnor}@ec-lyon.fr

[‡]STMicroelectronics (Canada) Inc.
Ottawa, Canada
Email: Pierre.Paulin@st.com

Abstract— State-of-the-art System-on-Chip (SoC) consists of hundreds of processing elements, while trends in design of the next generation of SoC point to integration of thousand of processing elements, requiring high performance interconnect for high throughput communications. Optical on-chip interconnects are currently considered as one of the most promising paradigms for the design of such next generation Multi-Processors System on Chip (MPSoC). They enable significantly increased bandwidth, increased immunity to electromagnetic noise, decreased latency, and decreased power. Therefore, defining new architectures taking advantage of optical interconnects represents today a key issue for MPSoC designers. Moreover, new design methodologies, considering the design constraints specific to these architectures are mandatory. In this paper, we present a contention-free new architecture based on optical network on chip, called Optical Ring Network-on-Chip (ORNoC). We also show that our network scales well with both large 2D and 3D architectures. For the efficient design, we propose automatic wavelength/waveguide assignment and demonstrate that the proposed architecture is capable of connecting 1296 nodes with only 102 waveguides and 64 wavelengths per waveguide.

I. INTRODUCTION

The latest edition of ITRS (International Technology Roadmap for Semiconductors) [7] emphasizes "More Than Moore's Law" trend. This trend focuses on system integration rather than transistor density, allowing for both functional and technological diversification in integrated systems. The functional diversification allows for non-digital functionalities to migrate from the board level into chip-level. This allows for integration of new technologies that enable high performance, low power, high reliability, low cost, high design productivity.

Moreover, technology scaling down to ultra deep submicron domain provides for billions of transistors which enable the integration of hundreds of cores on a single chip. These cores, running at a high clock frequency, create a need for high data bandwidth and increased parallelism. Therefore, the role of interconnect becomes a dominant factor in performance. Designing such systems using traditional electrical interconnect poses a significant challenge: due to capacitive and inductive coupling [6] interconnect noise and propagation delay of global interconnect increase. Increase in propagation delay requires global interconnect to be clocked at a very low rate, which puts limits to achievable bandwidth and overall system performance. Some attempts were made to solve this problem using different interconnect architectures, such as using insertion of repeaters on interconnect [1] or using pipelined global interconnect. However, use of pipelining

leads to higher data transfer delays, due to the large number of pipeline registers required, and to increase in power consumption, both due to the number of additional registers and due to the increased operating frequency. Therefore, a new on-chip interconnect technology that can overcome the problems of electrical interconnect is highly desirable.

Use of Optical Network-on-Chip (ONoC) promises to deliver significantly increased bandwidth, increased immunity to electromagnetic noise, decreased latency, and decreased power. Aside of physical properties, use of wavelength routing and Wavelength Division Multiplexing (WDM) [18] contributes to the valuable properties of optical interconnect. For traditional routing, a part of the message contains the destination address, while for wavelength routing the chosen wavelength determines the destination address, therefore enabling low contention or even contention-free routing. WDM allows for multiple signals to be transmitted simultaneously, facilitating higher throughput. The current technology is mature enough to allow this integration, thanks to CMOS-compatible optical components, such as light sources [11], waveguides [12], modulators [15], [16], and detectors [17], [19].

Defining new architectures taking advantage of optical interconnects represents today a key issue for MPSoC designers. This paradigm shift requires new methodologies for the efficient design. The design methodologies have to take into account the new constraints specific to optical interconnect. For instance, the number of waveguides and wavelengths used for a design is limited for feasibility, variability, power consumption, area and cost considerations.

In this paper we propose a novel architecture, Optical Ring Network-on-Chip (ORNoC). ORNoC is contention-free (no need for arbitration) network with high throughput and low latency. In this architecture, a wavelength is reused for multiple communications on a single waveguide, to take into account the design constraints. Consequently, fewer waveguides are required and the scalability is facilitated. Contributions of this paper are twofold: i) novel contention-free architecture providing wavelength sharing and ii) methodology for wavelength/waveguide assignment for efficient design.

The paper is organized as follows. Section II discusses related work. Section III describes the proposed architecture and Section IV presents use scenarios. Section V presents the wavelength/waveguide assignment design methodology and Section VI presents the experimental results. Section VII concludes the paper.

II. RELATED WORK

Several contributions address ONoC design exploiting both electrical and optical NoC technologies. An approach using electrical interconnects for control flow and optical interconnects for data flow was proposed in [24]. The electrical signal precedes the optical one in order to reserve the optical path. Therefore, optical communications may be delayed until an optical path becomes free, resulting in contention delay. Therefore, this type of network is not contention-free. The same observation is made for the fat tree and the mesh ONoC proposed in [5], [4]. In [22], [10], electrical interconnects manage local communication while an optical interconnect is responsible for global communications. However, such Single-Write-Multiple-Read (SWMR) implementation implies that each wavelength flowing through the ONoC must be assigned to a given optical network interface, avoiding parallel communications through the same wavelength. Such technique drastically affects the ONoC scalability. Because we use WDM and reuse wavelengths to realize more than one communication on the same waveguide at the same time, ORNoC does not suffer from this inconvenient. The Corona architecture [25], [2] follows a Multi-Write-Single-Read implementation that requires arbitration to manage write conflicts. Arbitration is not required in ORNoC. The Firefly architecture [21] extends the prior work by proposing the implementation of reservation-assisted SWMR buses. The main objective is to reduce the power consumption of optical communications by using an initialization packet in charge of turning-on data receiver resources. As a drawback, extra-latency is required compared to the SWMR technique and the network throughput rapidly decreases with the token round-trip latency [20]. FlexiShare architecture [20] reduces this drawback by allowing the injection of a new token each cycle (token stream arbitration). In contrast, ORNoC does not require arbitration. Only [8] and [3] consider contention-free ONoC, but they do not consider any method to reduce the implementation complexity when total connectivity is not required. The cores of an ATAC processor [13], [23] are connected via an electrical and an optical network. The optical network is used for global broadcasting. Its topology is most similar to the proposed ORNoC, however, the contention-free property is based only on WDM, while in ORNoC it is based on both WDM and wavelengths reuse. Moreover, ATAC does not support simultaneous communications between one source and multiple destinations, unless it is broadcast of the same message. Our approach has the potential for fewer waveguides/wavelengths.

III. SYSTEM AND ORNoC ARCHITECTURE

This section presents the overall architecture proposed in this work and the ORNoC network. We distinguish the optical portion from the electrical portion of the architecture. The electrical portion is composed of a set of computing nodes interconnected through a NoC while the optical part integrates the ORNoC network (i.e. including on-chip lasers, waveguides, etc.). We assume that nodes in electrical portion are grouped into clusters, in 2D scenario, and into electrical layers, in 3D scenario. In both scenarios, the optical portion is placed on a dedicated layer. Two types of communications are distinguished:

- *intra-cluster/intra-layer* communications are used for data transfers between the nodes situated within the same electrical cluster/layer
- *inter-cluster/inter-layer* communications are used for data transfers between nodes of different electrical clusters/layers.

The ORNoC is dedicated to inter-cluster/inter-layer type of communications. This communication requires a routing composed of

three main steps: (1) the electrical routing from the source node to appropriate Optical Network Interface (ONI), (2) the optical routing and (3) the electrical routing from the ONI to the destination node. The inter-cluster/inter-layer transceiver communication starts when an ONI receives a data and a destination ID. The data is serialized and the appropriate CMOS driver circuit then modulates the current flowing through the microresonator. The intensity of light emission is modulated according to the data bit values, achieving the electro-optical conversion. The signal enters into the ORNoC, gets routed inside it, and is finally received by a receiver's ONI. In the receiver ONI, the photodetector starts the opto-electronic conversion by converting a flow of photons into a photocurrent. An analog signal gets transmitted to the CMOS receiver circuit that converts the analog signal to a digital one, which is then deserialized. Data is finally injected into the electrical NoC and it is transmitted to the destination node.

In the following sections, we present details of ORNoC, where all of ONI in different clusters are able to communicate with each other (i.e. full connectivity).

A. Global View

The principle of ORNoC is illustrated using an example shown in Figure 1. Figure 1(a) shows 8 ONIs (A, B, etc.) connected with a single waveguide. The unmatched feature of ORNoC is that the same wavelength may be used on a single waveguide to concurrently realize multiple communications, as illustrated in Figure 1(b). In this representation, the single (physical) waveguide is represented as multiple (virtual) rings, each one being associated to a given wavelength. In the figure, 6 rings are represented, meaning that 6 different wavelengths can be used. For example, we may use λ_1 , shown in the smallest ring in the Figure 1(b), to realize communications: from ONI_A to ONI_B , from ONI_B to ONI_C , etc. When considering only λ_1 , the portion p_1 of the waveguide is used exclusively for ONI_A to ONI_B communication. However, because of WDM, many signals with other wavelengths may be used to realize other communications on the same portion of the waveguide. In the shown example, λ_3 , shown in blue on the second ring, is used to realize communication between ONI_A and ONI_D .

Fig. 1. Example of ORNoC highlighting multiplexed wavelengths implementing connections.

In general, the use of WDM allows multiple communications in the waveguide, at the same time. However, due to the technological constraints there is an upper bound for the number of wavelengths

that may be multiplexed. Therefore, we may use multiple waveguides and different directions of rotation to provide sufficient number of communication channels while conforming to the technological restrictions that limit the number of wavelength and waveguides.

B. Electrical Portion of ONI

Aside of the serializing data and modulating the signal, electrical portion of ONI transmitter is in charge with selecting an appropriate wavelength-waveguide pair to be used. We propose an architecture where a wavelength-waveguide pair is statically determined based on the destination. The details of this mapping are shown in Section V.

On the transmitter side, after analog-to-digital signal conversion and deserialization, the data is stored in a buffer that is dedicated to the particular wavelength. The buffers might become a bottleneck of the system. A simple solution for the buffer overflow is to have the overflow detected and an appropriate message sent to the sender. This can be done at a protocol level and will be investigated in future works.

C. Optical Portion of ONI

The optical portion of ONI is shown in Figure 2. ONI's receiver and transmitter parts consist of sets of n_r and n_t microresonators, respectively, each set characterized by its resonance wavelengths: λ_{ri} , for receiver and λ_{tj} , for transmitter part. In the receiver part, as shown in the figure, trajectory of the signal depends on the value of the signal's wavelength λ_s . Therefore, we define two modes of operation on receiver part:

- *ejection*: when $\lambda_s = \lambda_{ri}$ ($i = 0, \dots, n_r - 1$), the signal will couple into the microresonator, with the same wavelength as the signal, and then couple out into the perpendicular waveguide (e.g. signals with λ_b and λ_g shown in blue and green in the figure).
- *pass through*: when $\lambda_s \neq \lambda_{ri}$ ($i = 0, \dots, n_r - 1$), the signal will propagate along the same (ring's) waveguide. The red signal (i.e. λ_r) illustrates this operation.

Transmitter part adds one more state to ONI's functionality:

- *inject*: once data reaches ONI, a driver and modulator couple in the signal into an appropriate microresonator, which then couples it out to the waveguide (e.g. signals with λ_g and λ_p , shown in green and purple in the figure).

Fig. 2. Microresonators in optical portion of the ONI.

The proposed ONI architecture allows reuse of the same wavelength to realize multiple communications in the same waveguide, at the same time. For instance, in the example shown in Figure 1, the same wavelength may be used to realize the communication $ONI_A \rightarrow ONI_B$, as well as the communication $ONI_B \rightarrow ONI_C$, etc. The wavelength reuse translates into smaller overall number of wavelengths used in the system and, therefore in better scalability.

D. Use of Multiple Rings

The ORNoC contention-free property is made possible through the use of multiple wavelengths, both in WDM and wavelength reuse context. However, the number of wavelengths is currently technologically limited (e.g. 16), which may restrict the architecture scalability. In order to support large scale architectures, ORNoC supports the use of multiple rings. Each ring is independent, therefore the same set of wavelengths can be reused for each ring. The layout for the multi-ring topology allows for lower power losses, because there are no waveguide crossings. Moreover, our network has statically determined paths, so we can fine tune the on-chip laser output power to fit the requirements, while in Corona [25], [2] a worst case output power needs to be used for all the off-chip laser sources.

IV. ORNoC USE SCENARIOS

In this section, we present a 2D typical ORNoC use scenario and a more prospective 3D one.

A. 2D Context

In the 2D context, we propose a hierarchical network topology, similar to Firefly architecture [21], where cores are grouped in clusters. The cores within the cluster are connected using electrical network and they communicate with each other via the electrical network (intra-cluster communication). The nodes from different clusters are connected with ORNoC (inter-cluster communication). The layout of the proposed architecture for 3×3 clusters configuration is shown in Figure 3(b). The layout is fairly regular; this allows for automatic place-and-route and straightforward loss estimation. Moreover, the waveguide forming ring does not have any crossings, has relatively small number of bends connected with otherwise straight-line waveguides.

Assuming that each cluster is connected to the ORNoC through a single ONI (i.e. for 3×3 clusters architecture, 9 ONIs are used), we define a *Connectivity Matrix* (CM). The Figure 3(b) illustrates the connectivity values for the CM of the system shown in Figure 3(a). Rows and columns of CM contain source ONIs and target ONIs, respectively. If the communication is possible the value is '1', otherwise it is '0'. Since a full connectivity scenario is considered and self communications (e.g. ONI_A to ONI_A) are avoided, only the top-left to bottom-right diagonal contains '0' values.

Fig. 3. 2D architecture.

B. 3D Context

Figure 4(a) illustrates the proposed 3D architecture integrating ORNoC. It is composed of a set of stacked electrical layers and one optical layer. All electrical layers are connected to the optical layer using electrical vertical TSVs (Through Silicon Vias [14])

that upload and download the data between electrical and optical layers. Similarly to the 2D scenario, the cores in electrical layers are interconnected with electrical NoC while ORNoC is used for inter-layer communications. However, we also assume that each layer may have multiple accesses to the ORNoC.

As an example, we consider a 3D architecture with 2 layers and 4 (i.e. 2×2) ONIs per layer. The resulting CM is represented in Figure 4(b). ONIs located on the same layer never communicate with each other using the ORNoC (value '0' in CM).

Fig. 4. 3D architecture.

V. METHODOLOGY FOR DESIGNING EFFICIENT ORNoC

In order to design an efficient and feasible ORNoC we propose a design methodology that takes into account design constraints and design/connectivity requirements.

A. Design Methodology: Algorithm

The objective of proposed design methodology is to customize a generic ORNoC architecture for a given scenario by providing sufficient number of rings (i.e. waveguides) for supporting all the communications in a contention-free manner. Additionally, the number of rings must be minimized for power consumption and layout constraints purposes. For this purpose, we need to efficiently use available wavelengths on each portion of a ring.

We propose an algorithm that performs wavelength assignment and builds a custom ORNoC structure according to given requirements:

The inputs of the algorithm are: the maximum number of wavelengths per ring (max_wl), the number of ONI per cluster ($oni_per_cluster$), and the number of layers ($layers$). In 2D scenario, the number of layer equals one, where in 3D scenario, we consider having a single cluster on each layer, and therefore the number of ONIs on it is the value of $oni_per_cluster$. An element in the CM represents a communication between source-target couple. Each element has five fields: *connectivity*, *ring*, *portion*, *wl* and *processed*. The *connectivity* field equals '1' if corresponding source and target communicate via ORNoC (these values were discussed in Section IV), *ring*, *portion* and *wl* point to ring, portion of the ring and wavelength that are used to realize communication between source and target, while *processed* denotes that *ring*, *portion* and *wl* values have been assigned to the element.

Function *populate_connectivity_values* deduces value of *connectivity* field: in 2D scenario, we assume full connectivity, while in 3D scenario, we assume that only inter-layer communications are realized via ORNoC. Results of these assignments are shown in Figure 3(b) and Figure 4(b). Note that any other communication scenario may be handled with our methodology, which would require that the *connectivity* values are provided by a designer. Once the *connectivity*

Algorithm 1: Generate ORNoC

Data: $max_wl, oni_per_cluster, layers$

Result: configured ORNoC

```

1  $CM \leftarrow populate\_connectivity\_values(max\_wl, oni\_per\_cluster);$ 
2 while exists  $CM.element.processed = false$  do
3    $ring \leftarrow create\_ring(max\_wl);$  // alternate clockwise and
   counter-clockwise rotation
4   forall the unprocessed  $CM.element$  do
5     if  $CM.element.connectivity = 1$  then
6        $p[] \leftarrow ring.get\_portions(CM.element);$ 
7       forall the wavelength  $wl$  do
8         if  $wl$  available on all the portions then
9            $CM.element.ring \leftarrow ring;$ 
10           $CM.element.portion[] \leftarrow p[];$ 
11           $CM.element.wl \leftarrow wl;$ 
12           $CM.element.processed \leftarrow true;$ 
13          forall the portions of the communication path do
14             $p[i].wl.used \leftarrow true;$  // mark the
            wavelength  $wl$  as no more available
            on the set of ring portions
15          end
16          break;
17        end
18      end
19    else
20       $CM.element.processed \leftarrow true;$ 
21    end
22  end
23 end

```

values are obtained, the algorithm iterates in order to associate a communication path (i.e. a ring, a portion and a wavelength) to each element in CM (lines 2-23). Function *ring.get_portions* starts by considering use of only a single ring: we first consider short-range communications (ONI located close to each other on a ring) then longer-range ones. Once we identify the path (i.e. a set of ring portions $p[]$), we search for an available wavelength (lines 7-8). If the search finds an available wavelength, the wavelength is reserved, and the ring, portion and the wavelength values of $CM.element$ are populated. The element is also marked as processed. Once all the elements have been evaluated for a given ring, we check if another ring is necessary: if not, the process stops, otherwise a ring is added and all the non-processed elements are reconsidered. Note that, for multiple rings, we alternate the use of clockwise and counter-clockwise rotation ring in order to minimize the communication length, and thus the power consumption.

B. Design Methodology: Example

This section presents an example of wavelength assignment for a 3D architecture with 2 layers and 4 (i.e. 2×2) ONIs per layer (Figure 4(a)). The CM for this scenario is given in Figure 4(b). We also consider a maximum of 6 wavelengths per waveguide. Our algorithm is able to provide all required communications using only 2 waveguides. Figure 5(a) and 5(c) show the virtual view of waveguides where the communications are realized in clockwise and counter-clockwise directions, respectively. The waveguide illustrated in Figure 5(a) was firstly used to manage the communications; as a consequence, its ring portions are more intensively used compared to the second waveguides (in the latter, the largest circle is empty, thus a wavelength is not used). Figure 5(b) and 5(d) represent corresponding wavelength assignment portion on CM for two waveguides, showing that all the communications are correctly managed. The value of each node in CM gives the wavelength (i.e. the circle in the virtual representation) used to perform the communication.

Fig. 5. ORNoC with double ring a-b) virtual view for clockwise rotation waveguide and corresponding wavelength assignment c-d) virtual view of second waveguide with counter-clockwise rotation and corresponding waveguide assignment.

VI. CASE STUDY

We evaluated the efficiency of the ORNoC architecture and implemented the methodology for its automated design. We study the required number of waveguides as a function of the number of nodes, for different values of maximum number of wavelengths per waveguide. We use values 8, 10, 16 and 24 for the maximum number of wavelengths per waveguide in all the experiments. This choice is conform to the existing technology: for 5nm lithography and 90nm technology, for wavelengths between 1557-1583nm and microresonator radii from 1.0-2.5 μ m, waveguide with 16 multiplexed wavelengths have been fabricated [18], [26], [9]. This value is dependent on free spectral range (spacing between consecutive resonant frequencies: 50 nm), quality factor of a microresonator (500-800) and variability. It is projected that in near future, 24 wavelengths would be realistic implementation option.

A. 2D Context

Figure 6(a) shows the evolution of the number of required waveguides for 2D use scenario. We vary the number of clusters for this architecture: 4, 9, 16, 25 and 36. Please note that 4 clusters represent 2 $\times$ 2 scenario, 9 are 3 $\times$ 3, etc. The number of waveguides, as expected, increases with the increase of the number of clusters. The rate of increase of the number of wavelengths is biggest for the wavelength constraint that equals 8, and reaches 66 for 36 clusters. It worth to note that state-of-the-art SoC have up to hundred of cores, most of the proposed architectures suggest using 64 wavelengths and order of hundred of waveguides [25], [21], [8]. Indeed, the higher the number of tolerated wavelengths; the more communication paths can be multiplexed using the WDM technique.

Even though exact comparison is not possible, due to the different assumptions about how the networks operate, let us mention some

Fig. 6. Number of waveguides for different architecture configurations.

of the previously proposed architectures, alongside ORNoC. For connecting 64 (8 $\times$ 8) clusters Clos [8] optical network uses 64 waveguides, with total of 128 wavelengths. Corona [2] architecture is able to connect 256 (4 $\times$ 64) cores, using 4*64 waveguides for the crossbar and additional 1 waveguide for broadcast, with 64 wavelengths. For the same configuration, 4 $\times$ 64 and a constraint of 64 wavelengths, ORNoC requires only 26 waveguides.

B. 3D Context

We present the evolution of the number of waveguides as a function of the number of nodes for different maximum number of wavelengths in case of 2 and 4 electrical layers architecture.

Figure 6(b) and 6(c) show the number of waveguides required for ORNoC using 2 and 4 layers, respectively, for different number of ONIs per layer: 4, 9, 16, 25 and 36. In both cases, the number of waveguides increases with the number of layers and the number of nodes per layer. The rate of increase is faster for more stringent wavelength constraint: for maximum of 8 wavelengths per waveguide, and for 36 ONIs per each layer, ORNoC requires 132 and 794

waveguides, for 2 and 4 layers, respectively. This is the total number of waveguides required for partial connectivity between 72 (2 layers) or 144 (4 layers) ONI. In case of more realistic scenario, where a maximum of 16 wavelengths is allowed, only 66 and 398 waveguides are required, while for 24 wavelengths, the number of waveguides drops further to 44 and 264, for 2 and 4 layers, respectively.

Furthermore, if we consider 64 wavelengths allowed per waveguide, such as used in Corona and ATAC architecture, in order to connect 36 nodes per layer, the required number of waveguide would be 18 and 102, for 2 and 4 layer architectures, respectively. Another common design choice would be to form multiple clusters on a layer. If we consider 9 nodes per cluster and 36 clusters per layer, total of 1296 nodes on 4 layers could be connected with only 398, 264 and 102 waveguides, in case of 16, 24 and 64 wavelengths, respectively. Therefore, we may conclude that ORNoC allows for small number of waveguides, and it also shows potential for even further reduction in the number of waveguides with the advances of the optical on-chip technology.

VII. CONCLUSION

We propose ORNoC, a contention-free Optical NoC and its design methodology. The unmatched feature of ORNoC is that the same wavelength can be used to realize multiple communications on the same waveguide, at the same time, with no arbitration required. The proposed design methodology implements automatic wavelength-/waveguide assignment. We show that our network scales well with both large 2D and 3D architectures and our design methodology efficiently exploits the novel ORNoC architecture. Using ORNoC architecture and its design methodology we efficiently connect not only hundreds of processing elements, as in state-of-the-art SoC, but also thousand of processing elements, that are predicted for the near future. Our experiments show that the proposed architecture is capable of connecting 1296 nodes with only 102 waveguides, in case of 64 wavelengths per waveguide, making it valuable asset for current and future SoC designs with optical interconnect. Our future work includes power loss modeling and estimation for the ORNoC and investigate influence of errors in optical interconnect on overall system reliability.

REFERENCES

- [1] V. Adler and E.G. Friedman. Repeater Design to Reduce Delay and Power in Resistive Interconnect. *IEEE Transactions on Circuits and Systems II: Analog and Digital Signal Processing*, 45(5):607–616, May 1998.
- [2] R. G. Beausoleil, J. Ahn, N. Binkert, A. Davis, D. Fattal, M. Fiorentino, N. P. Jouppi, M. McLaren, C. M. Santori, R. S. Schreiber, S. M. Spillane, D. Vantrease, and Q. Xu. A Nanophotonic Interconnect for High-Performance Many-Core Computation. In *Proceedings of the 16th IEEE Symposium on High Performance Interconnects*, pages 182–189, 2008.
- [3] Mark J. Cianchetti, Joseph C. Kerekes, and David H. Albonesi. Phastlane: a Rapid Transit Optical Routing Network. In *Proceedings of the 36th International Symposium on Computer Architecture*, ISCA, pages 441–450, 2009.
- [4] Huaxi Gu, Jiang Xu, and Zheng Wang. A Novel Optical Mesh Network-on-Chip for Gigascale Systems-on-Chip. In *Proceedings of the IEEE Asia Pacific Conference on Circuits and Systems*, pages 1728–1731, December 2008.
- [5] Huaxi Gu, Jiang Xu, and Wei Zhang. A Low-Power fat Tree-Based Optical Network-on-Chip for Multiprocessor System-on-Chip. In *Proceedings of the conference on Design, Automation and Test in Europe*, DATE, pages 3–8, 2009.
- [6] R. Ho, K.W. Mai, and M.A. Horowitz. The Future of Wires. *Proceedings of the IEEE*, 89(4):490–504, April 2001.
- [7] International technology roadmap for semiconductors. <http://www.itrs.net/>.
- [8] Ajay Joshi, Christopher Batten, Yong-Jin Kwon, Scott Beamer, Imran Shamim, Krste Asanovic, and Vladimir Stojanovic. Silicon-Photonic Clos Networks for Global on-Chip Communication. In *Proceedings of the 3rd ACM/IEEE International Symposium on Networks-on-Chip*, NOCS, pages 124–133, 2009.
- [9] A. Kazmierczak, M. Briere, E. Drouard, P. Bontoux, P. Rojo-Romeo, I. O'Connor, F. Letartre, X. Gaffiot, R. Orobtcouk, and T. Benyattou. Design, Simulation, and Characterization of a Passive Optical Add-Drop Filter in Silicon-on-Insulator Technology. *IEEE Photonics Technology Letters*, 17(7):1447–1449, 2005.
- [10] Nevin Kirman, Meyrem Kirman, Rajeev K. Dokania, Jose F. Martinez, Alyssa B. Apsel, Matthew A. Watkins, and David H. Albonesi. Leveraging Optical Technology in Future Bus-based Chip Multiprocessors. In *Proceedings of the 39th Annual IEEE/ACM International Symposium on Microarchitecture*, MICRO, pages 492–503, 2006.
- [11] M. J. Koblinsky. On-Chip Optical Interconnects. *Intel Technology Journal*, 08(02):129–141, October 2004.
- [12] S.J. Koester, G. Dehlinger, J.D. Schaub, J.O. Chu, Q.C. Ouyang, and A. Grill. Germanium-on-Insulator Photodetectors. In *2nd IEEE International Conference on Group IV Photonics*, pages 171–173, 2005.
- [13] Jifeng Liu, James Psota, Nathan Beckmann, Jason Miller, Jurgen Michel, Jonathan Eastep, George Kurian, Lionel Kimerling, Anant Agarwal, and Mark Beals. ATAC: A Manycore Processor with On-Chip Optical Network. Technical report, MIT-CSAIL-TR-2009-018, May 2009.
- [14] Igor Loi, Federico Angiolini, and Luca Benini. Supporting Vertical Links for 3D Networks-on-Chip: Toward an Automated Design and Analysis Flow. In *Proceedings of the 2nd international conference on Nano-Networks*, Nano-Net, pages 1–5, 2007.
- [15] Yehia Massoud, Naomi Halas, and Peter Nordlander. Subwavelength Nanophotonics for Future Interconnects and Architectures. Invited talk, NRI SWAN Center, Rice University, 2008.
- [16] D. Miller. Device Requirements for Optical Interconnects to Silicon Chips. *Proceedings of the IEEE*, 97(7):1166–1185, 2009.
- [17] Jacob R. Minz, Somaskanda Thyagaraja, and Sung Kyu Lim. Optical Routing for 3D System-on-Package. In *Proceedings of the conference on Design, Automation and Test in Europe*, DATE, pages 337–338, 2006.
- [18] I. O'Connor, F. Mieyeville, F. Gaffiot, A. Scandurra, and G. Nicolescu. Reduction Methods for Adapting Optical Network on Chip Topologies to Specific Routing Applications. In *Proceedings of the Design of Circuits and Integrated Systems*, DCIS, November 2008.
- [19] Ian O'Connor and Frédéric Gaffiot. On-Chip Optical Interconnect for Low-Power. In Enrico Macii, editor, *Ultra Low-Power Electronics and Design*, pages 21–39. Springer US, 2004.
- [20] Yan Pan, J. Kim, and G. Memik. FlexiShare: Channel Sharing for an Energy-Efficient Nanophotonic Crossbar. In *Proceedings of the 16th IEEE International Symposium on High Performance Computer Architecture*, HPCA, pages 1–12, January 2010.
- [21] Yan Pan, Prabhat Kumar, John Kim, Gokhan Memik, Yu Zhang, and Alok Choudhary. Firefly: Illuminating Future Network-on-Chip with Nanophotonics. In *Proceedings of the 36th annual International Symposium on Computer Architecture*, ISCA, pages 429–440, 2009.
- [22] Sudeep Pasricha and Nikil Dutt. ORB: an on-Chip Optical Ring Bus Communication Architecture for Multi-Processor Systems-on-Chip. In *Proceedings of Asia and South Pacific Design Automation Conference*, ASP-DAC, pages 789–794, 2008.
- [23] J. Psota, J. Miller, G. Kurian, H. Hoffman, N. Beckmann, J. Eastep, and A. Agarwal. ATAC: Improving Performance and Programmability With on-Chip Optical Networks. In *Proceedings of IEEE International Symposium on Circuits and Systems*, ISCAS, pages 3325–3328, 2010.
- [24] Assaf Shacham, Keren Bergman, and Luca P. Carloni. Photonic Networks-on-Chip for Future Generations of Chip Multiprocessors. *IEEE Transactions on Computers*, 57:1246–1260, September 2008.
- [25] Dana Vantrease, Robert Schreiber, Matteo Monchiero, Moray McLaren, Norman P. Jouppi, Marco Fiorentino, Al Davis, Nathan Binkert, Raymond G. Beausoleil, and Jung Ho Ahn. Corona: System Implications of Emerging Nanophotonic Technology. In *Proceedings of the 35th Annual International Symposium on Computer Architecture*, ISCA, pages 153–164, 2008.
- [26] L. Zhang, M. Yang, Y. Jiang, E. Regentova, and E. Lu. Generalized wavelength routed optical micronetwork in network-on-chip. In *Proceedings of the 18th IASTED international conference on Parallel and Distributed Computing and Systems*, PDCS, pages 698–703, November 2006.