Communication Aware Design Method for Optical Network-on-Chip

Johanna Sepúlveda[1,2], Sebastien Le Beux[2], Jiating Luo[1], Cedric Killian[1],
Daniel Chillet[1], Hui Li[2], Ian O'Connor[2], Olivier Sentieys[1]
[1]INRIA, IRISA, University of Rennes, Lannion 22300, France
[2]Lyon Institute of Nanotechnology, Ecole Centrale de Lyon, Ecully 69134, France
martha-johanna.sepulveda-florez@inria.fr

***Abstract*—Optical technology promises to solve the bottleneck communication in Multiprocessor Systems-on-Chip (MPSoCs) by integrating high speed interconnections. From the system point of view, one of the most critical problems in optical communication is channel bandwidth, strongly influencing the system performance and cost. Channel bandwidth establishes the number of waveguides and wavelengths to serve each communication request. Dynamic approaches allow to define the channel bandwidth at the runtime, which leads flexible systems but may turn their performance unpredictable. Design-time approaches appear as an attractive alternative to define the channel bandwidth for systems that require performance guarantees and simple solutions. A method for exploring the channel bandwidth design space is thus mandatory in order to identify the best communication channel size. In this work we explore the trade-off among channel bandwidth alternatives, performance, area and power. We show that the channel size has a strong impact on the system performance and cost. We employ synthetic and real application traffic which has been executed on Gem5. As a result we show that different channel bandwidth can improve the execution time of an application up to 75% while including low area and power penalties.**

Keywords-component; Silicon Photonics; Network-on-Chip; Interface; Design exploration

I. Introduction

Current applications promote the quest of semiconductor industry by efficient more-than-Moore devices. Such systems must be able to support complex functionalities under tight performance requirements and smaller size and low power constraints [1]. As a result, the amount of hardware components integrated into a single chip has increased significantly. Current systems are able to integrate hundreds of cores turning the inter-core communication the performance bottleneck of high performance Multi-Processors System-on-Chip (MPSoCs).

3D system integration appeared as a promising technology for solving the communication problem. 3D-MPSoCs are able to reduce the distance among the cores by stacking different hardware components mapped onto several layers of dies into a single chip. Intra-layer communication is performed by electrical networks (Networks-on-Chip; ENoCs), which uses routers and links to perform the data exchange. Inter-layer communication is carried out by buses of Through-Silicon-Vias (TSVs). Despite the alleviation of the communication problem, electrical technology is not enough for solving it. However, the 3D-technology incursion also allows the integration of different technologies into a single chip. Such characteristic has led the exploration of new communication alternatives.

Optical interconnection appears to be a promising disruptive technology to achieve high speed on-chip communications. It integrates microresonators and lasers to receive and transmit through waveguides data encoded as light. Compared to an electrical communication approach, the optical communication enhances the bandwidth density and energy efficiency, by packing dozens of wavelengths into a single waveguide. Each wavelength is able to operate from 5 to 10 Gb/s [2]. Optical Networks-on-Chip (ONoCs) are integrated into an MPSoC by stacking an optical layer atop the electrical layer. An optical interconnected MPSoC is shown in Fig. 1. The optical die integrates the microresonators, lasers and waveguides required for optical communication. The electrical layer includes the Intellectual Property (IP) components required to process and store the data.

Optical resources are limited and they must be shared among the different communication requests [3]. Channel bandwidth a critical design parameter since it considerably impacts system performances and cost. It dimensions the network interface by deciding the number of waveguides and the number of wavelengths that must be used for any communication request.

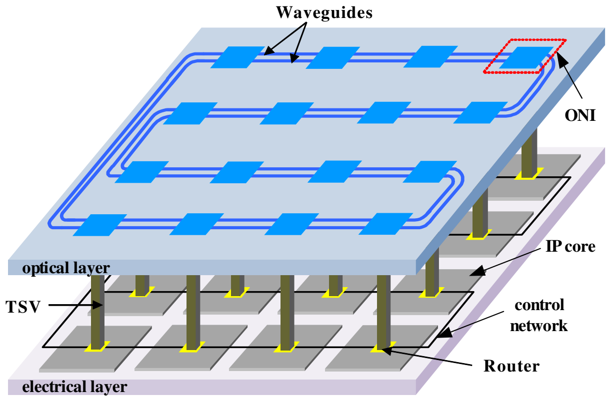

Fig. 1. General Optical NoC-based MPSoC architecture. Communication structure is highlighted.

Thus, it defines the number of laser, microresonators and waveguides integrated at the system. The channel must meet the performance requirements while minimizing the channel bandwidth allocation time (i.e. the communication latency) and the cost of the channel controller. The channel bandwidth, or the channel size, is defined by the number of optical resources involved in each communication. The communication channel size can be decided at design time, during the optical MPSoC design stage, or at runtime, during the system operation. Each time a request is performed, the allocator is able to reserve, among the available optical communication resources, the channel size defined at design-time or runtime. The runtime approaches reserve for each request different numbers of waveguides and wavelengths in order to satisfy the performance requirements. Despite the flexibility, runtime approaches are characterized by costly allocator modules and long allocation time [3][4]. For applications with tight performance requirements, fixing the communication channel size at design-time turns an attractive alternative. Thus, a design space exploration of the optical network interface is mandatory in order to find the best communication channel size. Coarse-grained communications (high bandwidth communications) may be advantageous for reducing transmission time, but it may increase the channel allocation time.

Several design parameters of optical interconnections have been explored in previous works. However, the impact of the communication channel size on the system performance and cost has not been addressed. This work addresses the design space exploration of optical interfaces in order to define the best communication channel size at design-time. We show that the channel size has a strong impact on the system performance and cost. Our approach leads to enhancements of performance up to 75% and reduction of area and power up to 55% and 23%, respectively.

The remaining text is organized as follow. Section II presents an overview of previous works. Section III describes the optical MPSoC architecture and the channel bandwidth problem. Section IV presents the communication channel size exploration. Experiments and results are reported in Section V. Finally conclusions are provided in Section VI.

II. RELATED WORK

Many of the previous research works in the area of optical interconnections are focused on improving performance and reducing costs by exploring different topologies and device configurations. However, it has been shown in [3][4] that channel bandwidth is an important parameter on optical communication design. According to the stage where the communication channel size is defined, previous works can be classified into design-time or runtime approaches. Table I. summarizes the characteristics of the previous works.

The definition of the channel bandwidth at runtime is a technique employed in [4][5][6][7]. These works use a controller which selects the number of waveguides and wavelengths based on an allocation algorithm.

TABLE I. RELATED WORK

Ref.	Laser	Definition of the channel bandwidth	Allocation Waveguide/ Wavelength	Exploration channel size
[4]	On-chip	Runtime	wg/wl	No
[5]	On-chip	Runtime	wg/1	No
[6]	Off-chip	Runtime	1/1	No
[7]	Off-chip	Runtime	1/16	No
[8]	Off-chip	Design-time	1/1	No
[9]	Off-chip	Design-time	1/1	No
[18]	On-chip	Design-time	1/1	No
Ours	On-chip	Design-time	wg/wl	Yes

Despite of the good results, the algorithm complexity is high and leads huge area and performance overhead. Moreover, the performance predictability is not guaranteed.

In [7], a reconfiguration algorithm is used to improve performance in terms of throughput and latency by adapting the available network bandwidth to the application demand. However, this approach highly depends on traffic monitoring and the fully reconfigurable design leads to area overhead and higher complexity.

The design-time approach is used by the works of [8][9][18]. These works reserve a wavelength for dedicated communications among a pair of components. This approach guarantees the communication for every request. However, it is not scalable and it may lead to underutilization of optical resources. The work of [18] shows the strong impact of the configuration parameters of the electro-optical interface on the overall system performance. This work explores the parameters: buffer size, flit width and modulation speed of a wavelength-routed ONoC interface. Each interface is composed by symmetric transmission and reception parts. At the transmission part, each initiator has an output for each possible target. Targets outputs are fed through a serialized dual clock FIFO (DC FIFO). At the reception part, each target has an input. Received data is deserialized and stored in a DC FIFO. By synthesizing each ONoC transmission and reception parts into 40 nm technology, the authors determine the latency, power and area of each component. These values show that the DC FIFOs are the components with largest latency in the interface (3.31 ns and 3.81 ns for control and data FIFOs at the transmission part and 2.91ns and 2.68 at the reception part). When simulating a 4x4 ONoC under different traffic scenarios, they show that most of the latency of the ONoC is due the interface.

None of the previous works explored the communication channel size. This parameter has a huge impact on the interface design. In this work we present for the first time a study that evaluates the impact of the communication channel size on the performance of the system.

III. ARCHITECTURE

In this section, we first present an overview of the electro-optical architecture of the optical interconnected MPSoC. Then, the optical part of the network interface and the channel bandwidth problem are defined. Finally, the electrical part of the network interface is presented.

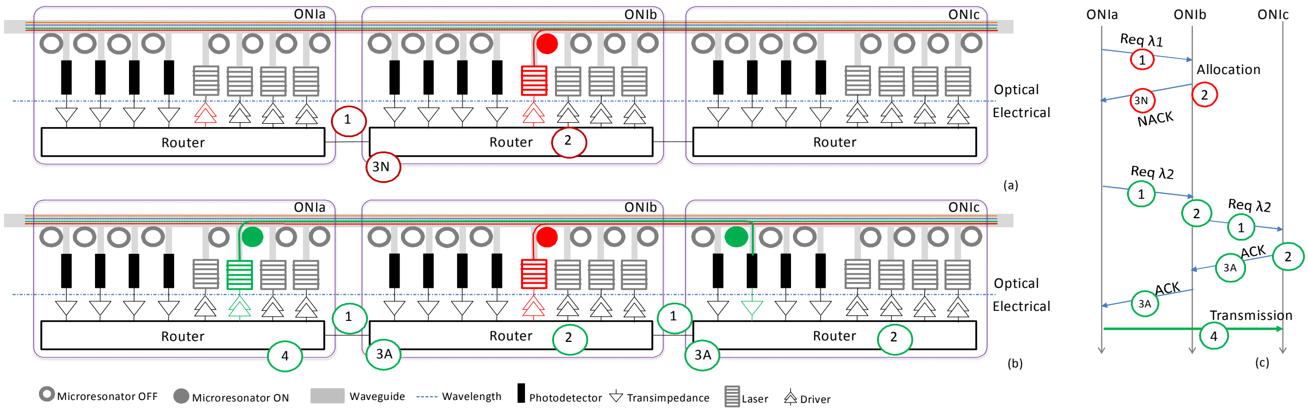

Fig. 2. Communication channel allocation example in ONoC . (a) Failed allocation attempt; (b) Succeeded allocation attempt; (c) protocol diagram.

A. *Overview*

Optical interconnected MPSoCs are composed of computation and storage IP components which exchange data through electro-optical networks. ENoCs integrate electrical routers and links to transmit the required messages for controlling the optical resources. ONoCs employ microresonators, lasers and waveguides to exchange the data information among the IPs through light signals. Electro-optical and optical-electrical converters made the interaction between ENoCs and ONoCs possible. In order to establish a communication between a *source IP* (which generates the request of optical communication) and a *target IP* (which is the destination of the communication), the reservation of the optical resources must be done. This is achieved by a channel allocator that aims to reserve a set of waveguides and wavelengths between the communicating IPs. Larger numbers of allocated waveguides and wavelengths provide higher channel bandwidth. However, it reduces the number of simultaneous communications in the ONoC due the limited and shared number of resources.

Integration of electrical and optical technologies into the same chip, is possible by three-dimensional integration. It stacks an optical layer atop an electrical layer into a single chip. The electrical layer groups all the computing and storage IP together with the ENoC. Fig. 1 presents the general architecture of the optical interconnected MPSoC. The electrical layer integrates 16 IP cores linked through a 16-router ring ONoC. The optical layer integrates the optical resources constituted by microresonators, lasers and waveguides. Fig. 1 presents a two waveguides system. Lasers are used to emit the light, which is transmitted by the waveguides and dropped by the microresonators. Optical and electrical devices are integrated into a structure called Optical Network Interface (ONI). The detail of the internal configuration of the ONI is shown at Fig. 2. The interconnection among different layers is performed through vertical buses based on Through-Silicon-Vias (TSVs), which are conductive nails that extend out the back-side of a thinned-down die [10].

Among several topologies, the ring has shown the best performance and lowest cost [4][11]. In this kind of topology each electrical router has four bidirectional ports for linking two neighbor routers, an IP and the optical components of the ONI. The next subsections describe the optical and electrical architectures.

B. *Optical structure of the ONI*

Wavelength-Division-Multiplexed (WDM) allows data transmission encoded as light signals. It allows multiplexing the data on different carrier wavelengths. Each wavelength is able to establish a channel between any pair of components. The optical architecture is composed by microresonators, on-chip lasers called Vertical-Cavity Surface Emitting Lasers (VCSELs) [12] and waveguides. Microresonators and lasers are integrated at the reception (Rx) and transmission (Tx) parts of the ONIs.

Microresonators are small silicon rings able to drop part of the spectrum around a resonance wavelength from a waveguide. Their diameter determines the resonance frequency λ_R. Light wavelength λ_i that matches the resonance frequency of the microresonators ($\lambda_i = \lambda_R$) will be deviated from the incoming waveguide to an alternative path. VCSELs are integrated lasers that emit light at different wavelengths [12]. It is a layered structure composed of two n-type and p-type doped mirrors. Each laser transforms the electrical current into light. The waveguide transports the light signal efficiently among the components of the chip. It confines light through a line-shaped path.

Fig. 2 shows a three-ONI system (ONI_a, ONI_b and ONI_c) communicated by one waveguide that integrates four wavelengths (λ_1, λ_2, λ_3 and λ_4). Each ONI is able to drop and transmit a light signal characterized by these wavelengths. They are constituted by eight microresonators divided into four for the reception part and four for the transmission part. The microresonators at the transmission part are integrated for coupling purposes. Such a phenomenon is widely presented in [13]. Fig. 2(a) illustrates the activation of the first laser of the ONI_b. In Fig. 2(b) an optical communication between ONI_a and ONI_c is illustrated. ONI_a activates the second laser, thus emitting light with wavelength λ_2. The light signal crosses through the ONI_b without being dropped since the λ_2 microresonator is turned off. At ONI_c, the communication destination, the λ_2 microresonator is activated and thus, the light is deviated from the main waveguide to the photodetector. It converts incident optical power into photocurrent. Optical-electrical conversion is achieved by the photodetector and transimpedance amplifier [2]. As a result, the optical signal is transformed into an electrical signal and it is ready to be injected into the electrical architecture.

C. *Channel bandwidth problem*

As shown in Fig. 2(c), data exchange is performed in the optical structure by four steps. At the first step (1), the *source IP* requests an optical communication. This request is managed according to the channel bandwidth policy, at step (2). It verifies the availability of the optical resources and grants the channel to the requested communication. Such optical resources remain reserved until an ACKnowledgement (ACK) signal is received. Any other communication request will not be able to use the reserved resources until they are released, at the end of the *source IP* – *target IP* communication. When the allocation fails (due the lack of free communication resources), a Non-ACKnowledgement (NACK) signal is sent back to the *source IP,* as shown in (3N). Otherwise, the steps (1) and (2) are repeated at each ONI until reaching the *target* ONI. When the allocation is successful in the target ONI, an ACK signal is sent back through the reserved resources (3A). Once the channel is reserved, the optical communication takes place at step (4).

Fig. 2 shows the communication among a *source IP* linked at the ONI_a and a *target IP* linked to the ONI_c. Fig. 2(a) presents the initial state of the optical resources. At ONI_b, the laser λ_1 is being used, preventing any other communication for using the path among ONI_b and ONI_c through λ_1. Steps (1) and (2) at ONI_a lead as a result the allocation of λ_1. In order to reach the *target* ONI, steps (1) and (2) are repeated at ONI_b. However, as λ_1 is reserved in ONI_b, a NACK is generated in (3N). Fig. 2(b) shows another attempt to establish an optical channel, this time the allocated channel is λ_2. Steps (1) and (2) are performed at ONI_b. As λ_2 is available, the laser resource is reserved, and the allocation process can be started at ONI_c by executing (1) and (2). ONI_c allocates λ_2 and the (3A) can be performed. When the ACK signal arrives to the *source* ONI, the optical communication can be performed. At the end of the transmission the reserved resources are released.

Fig. 2(c) shows a communication sequence diagram between ONI_a and ONI_c. The first attempt of communication is performed through light with wavelength λ_1. However, this channel is already used by ONI_b., i.e. the communication fails. The second attempt succeeded and the optical channel is established with wavelength λ_2. Any fail in the channel assignment may degrade the performance of the system through an increase of the latency. Thus channel bandwidth has a critical role in the optical communication design.

D. *Electrical structure of the ONI*

The electrical part of the ONI is divided into the *Rx* part, which includes electrical receiver circuits to adapt the electrical signal, and in the *Tx* part, which includes drivers to control the lasers. Both parts are controlled by the electrical ENoC, that implements the protocol that allows the data exchange by optical means.

The ENoC is composed by routers and links that exchange the data required to control the optical resources. Routers are used to commute the data from an input to an output. Neighbor routers communicate through bidirectional links. Routers integrates four main components: i) *input buffers*, which store the data that request the router by one of its input ports; ii) *arbitration and routing algorithm*, which grants the utilization of the crossbar switch to one of the input buffers and selects the router output port to be employed for redirecting the incoming data; iii) *channel allocator*, which keeps track of the status of the optical resources and reserves the wavelengths required to create the optical channel communication; and iv) *crossbar switch*, which links input to output ports of the router. The router microarchitecture is illustrated in Fig. 3.

The total commutation latency L_R of the four-stages router is determined by (1), where t_{br} is the time spent by the storage at the input buffers, t_a is the arbitration and the route calculation time, t_c is the time spent by the channel allocation process and t_s is the switching time.

$$L_R = t_{br} + t_a + t_c + t_s \tag{1}$$

A successful commutation takes place when the channel is allocated, which means there are available optical resources for transmitting the data. When the communication channel size has been decided at design-time, a simpler channel allocator module is generated compared to runtime approaches. These latter solutions require the execution of complex algorithms that may increase t_c and thus the total communication latency per router.

The number of waveguides and wavelengths per communication can be defined at design-time, and thus fixed in the device, or at runtime, by adapting the channel bandwidth according to the communication requirements.

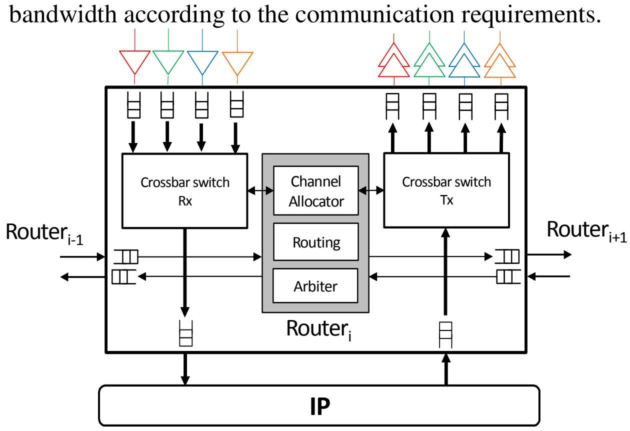

Fig. 3. Electrical router architecture.

However, despite the flexibility offered by the runtime alternative, the control complexity leads to prohibitive channel allocation latency for some sets of applications. Exploring the channel bandwidth design space is thus mandatory in order to identify the best communication channel size for some set of applications.

IV. Communication Channel Bandwidth

This section presents the design space exploration of the channel bandwidth and illustrates its purpose through illustrative examples.

A. Design space exploration

Applications executed on the optical MPSoC will define the communication pattern among the different components. The communication channel size may affect the execution time T_E of an application. It depends on the computation and communication capabilities of the system. T_E is given by (2), where t_p is related to the computation capabilities of the system. It is the processing time required to perform all the computing and storage operations at the IPs. Communication capabilities are related to t_R and t_o, where t_R is the time required to reserve all the resources in the light path, and t_o is the optical transmission time. Note that t_R, given by (3), corresponds to the sum of all the L_R of the routers (N_R) on the light path. As t_o is negligible, due the speed of the light, t_R becomes dominant. From (1), t_{br}, t_a and t_s are deterministic and depend on the electrical configuration of the routers. However, t_c depends on the communication channel size of the system. Therefore an exploration of this parameter results critical for the system performance.

$$T_E = t_p + t_R + t_o \tag{2}$$

$$t_R = \sum_{i=0}^{N_R} L_{R_i} \tag{3}$$

The channel bandwidth is defined by the couple (wg,wl), where wg represents the number of waveguides and wl the number of wavelengths reserved by each request.

Communication channel size defines different instances of the channel bandwidth. Each time a request is performed, the channel allocator will reserve (wg,wl). The importance of the channel size arises from the existence of limited and shared resources to perform the optical communication. Coarse-grained communication will guarantee high bandwidth communication, but compromise the allocation time ta, and thus t_R. The probability that all resources are available will depend on the total amount of channels, the number of channels that should be granted for each request and the number of requests. Fig. 4 shows the design space of the communication channel size expressed as a function of the couple (wg, wl). The interval of each axis varies from 1 to the maximum number of waveguides N_{WG} and maximum number of wavelengths N_{WL}.

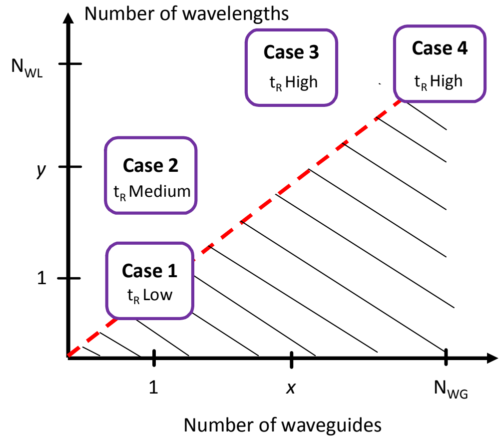

Fig. 4. Exploration space of the communication channel size.

B. Example of communication channel size

From the Fig. 4, four cases are identified as interesting points of exploration due to the efficient utilization of resources. *Case 1*, corresponds to the couple (1,1). It is characterized by the allocation of 1 wavelength per waveguide. That is, each time a request takes place, 1 waveguide is selected in which 1 wavelength must be reserved. For this case t_R is low, once the fine grained communication has higher probability to find free resources for building the lightpath. *Case 2*, corresponds to the couple (1,wl). It is defined by the allocation of wl wavelengths at 1 waveguide. For this case, t_R is medium, once the amount of resources required to establish the lightpath increases when compared to the *Case 1*. *Case 3*, corresponds to the couple (wg,N_{WL}). It allocates all the wavelengths of wg waveguides. For this case, t_R is high due the coarse grained communication. Finally, *Case 4*, corresponds to the duple (N_{WG},N_{WL}). It allocates all the optical communication resources per request. For this case, t_R is high.

In order to illustrate the four cases in an example, Fig. 5 represents a three-ONI system (ONI_a, ONI_b and ONI_c) similar to the one shown in Fig. 2. The difference is that we assume 4 waveguides in this example (N_{WG}=4). Each waveguide is able to integrate four wavelengths λ_1, λ_2, λ_3 and λ_4 (N_{WL}=4).

Three communication requests take place: i) $ONI_a \rightarrow ONI_c$; ii) $ONI_b \rightarrow ONI_c$; and iii) from ONI_c to a distant ONI not shown in the figure.

- The *Case 1* (1,1) is presented at Fig. 5(a), where the requests of ONI_a, ONI_b and ONI_c are allocated respectively at λ_1, λ_2 and λ_1 of the first waveguide. Note that although the ONI_a and ONI_c are allocated to the same wavelength λ_1, there is no data collision, once the light signal of ONI_a is dropped at the reception part of ONI_c. It precedes the transmission part of ONI_c. Therefore, where ONI_c injects the data, λ_1 is free.

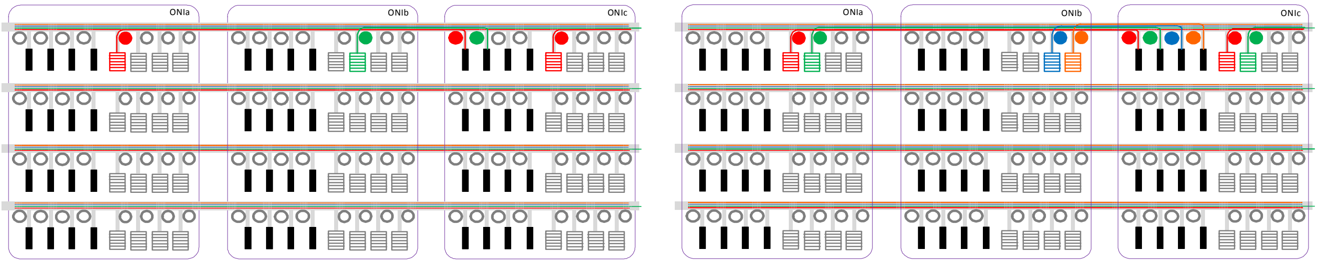

(a) Case 1 (1,1): 1 wavelength per waveguide.

(b) Case 2 (1,2): N wavelengths per waveguide (N=2).

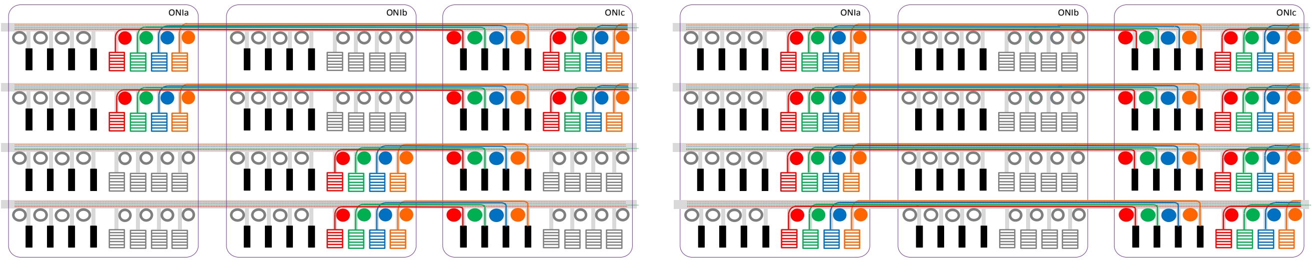

(c) Case 3 (2,4): All wavelengths per N waveguides (N=2).

(d) Case 4 (4,4): All wavelengths for all waveguides

Fig. 5. Different cases of communication granularities.

- The *Case 2* (1,2) is represented at Fig. 5(b). Each request is allocated into 2 wavelengths of the first waveguide. Therefore, the allocated wavelengths are λ_1 and λ_2 for ONI_a, λ_2 and λ_4 for ONI_b and λ_1 and λ_2 for ONI_c.
- The *Case 3* (4,2) is shown at Fig. 5(c). It allocates for each request all the wavelengths (N_{WL}=4) of 2 waveguides. That is, each time a request is performed, a total of 8 aggregated wavelengths are allocated. For the requests of ONI_a and ONI_c, the first and the second waveguides are used. The request of ONI_b is allocated at the third and fourth waveguides.
- *The Case 4* (4,4) is presented at Fig. 5(d). It allocates all the waveguides and wavelengths for each request (i.e. 16 aggregated wavelengths). Note that the request of ONI_a and ONI_c are served. However, the request of ONI_b fails. Therefore, the request of ONI_b must wait until ONI_a release the optical resources. This communication channel size causes an increase on the allocation time t_c of ONI_b and thus, in the total commutation latency L_R. However, the effect of this channel size on the execution time T_E, remains unknown, hence simulations are required. *Case 4* shows that bandwidth enhancements may result in higher allocation times.

V. Exploration Results

A. Model

In order to perform the design space exploration of the communication channel size, we have extended the Gem5 simulator [13]. Gem5 is an open source modular quasi-cycle accurate simulation environment which supports a wide variety of Instruction Set Architectures (ISAs) and all the components required for MPSoC simulation. By integrating C++ models of the optical interconnection components (ONI and waveguides) and modifying the *garnet* electrical interconnection network [14], we were able to model an optical interconnected MPSoC.

The channel allocation component was integrated on the electrical routers. Such characteristics allow the evaluation of the communication channel size impact on the system performance. Our platform is composed by a 8-tiles MPSoC interconnected by a ring-based optical interconnection. Each tile integrates an ARM processors (500MHz), 4-KB L1 cache and 32 bits as channel width. The system integrates 4 waveguides and 16 waveguides, each integrating 16 wavelengths. The channel size is varied from 1 to 256 aggregated wavelengths for the design space exploration purpose. The simulations take into account the allocation algorithm execution, which is given by the synthesis results from [17]. The modulation speed is varied from 500 Mb/s to 20Gb/s.

In order to test different channel allocation configurations, we employ synthetic and real application traffic. Synthetic traffic is composed by hot spot, uniform and transpose patterns, whose packet injection rate was fixed at 10%, 30% and 50%. The size of each data packet is 20 flits, where each flit is composed by 32 bits. Real traffic test was obtained after the execution of SPLASH-2 benchmark suite [16] at the platform.

B. Synthesis results

For obtaining the area and static power results of the channel allocator required for implementing the different allocation cases, we employ the synthesis results presented in [16]. This work employs a 28nm FDSOI technology, from Synopsys Design Vision environment (version H-2013.03-SP5). The results are summarized in Table II. Note that allocating all the resources (16,16) for any request results in simpler controllers. By changing the configuration from 2 (1,2) channels to 4 (1,4) channels, the savings in area and power are 30% and 9%, respectively. When comparing with the worst scenario, the savings can achieve up to 55% and 23%.

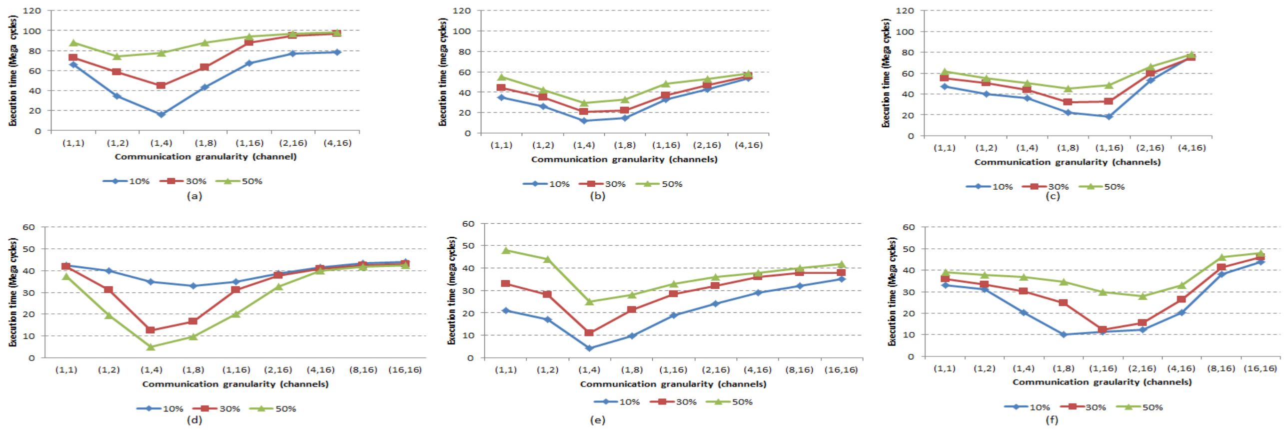

Fig. 6. Execution time for delivering 10000 packets under (a) Hot-spot for 4 waveguides; b) Uniform for 4 waveguides; c) Transpose for 4 waveguides; e)Hot-spot for 16 waveguides; f) Uniform for 16 waveguides; g) Transpose for 16 waveguides.

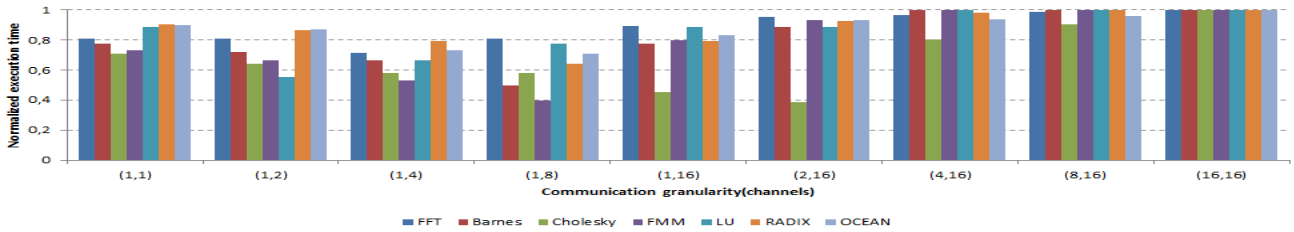

Fig. 7. Normalized execution time for Splash benchmarks

C. Execution time w.r.t. channel bandwidth

Fig. 6 gives the execution time for Hot-Spot, Uniform and Transpose traffic scenarios for various channel sizes defined by the couples (*wg,wl*). Results of Fig. 6(a) to (c) correspond to the 4 waveguides configuration and the Fig. 6(d) to (f) correspond to the 16 waveguides configuration. The huge impact of the channel size over the performance of the system can be observed. Each application presents an execution time that varies for each channel size. One of the most sensitive traffic is the hot spot, presented at Fig. 6(a) and (d), which presents curves characterized by higher slops. By changing the channel size from 2 (1,2) to 4 (1,4) channels in the case of 4 waveguides system, the enhancements of the execution time of the application for 10% and 30% injection rate are 53% and 23%, respectively. However, for the 50% of injection rate there was an increase of 5%. This value is result of the early saturation of the system caused by the intense traffic. It turns this traffic less sensible to the channel allocation parameter. For the case of 16 waveguides system, there were always enhancements for all the injection rates with the values of 75%, 59% and 12%, respectively.

Uniform traffic execution time results for 4 and 16 waveguides are presented in Fig. 6(b) and (e). The best results are obtained for 4 (1,4) channels for all the injection rates. For the 4 waveguides system, by modifying the channel allocation from 2 channels to 4 channels, the enhancements achieved for 10%, 30% and 50% injection rates are 54%, 40% and 30%, respectively. For the 16 waveguides system, the enhancements achieved are 75%, 60% and 43%, respectively.

For the transpose traffic, the best channel bandwidth highly depends on the traffic injection rate. This type of traffic highly stresses some paths inside the system, causing its early saturation. The best solutions varied from 8 (1,8) to 16 (1,16) channels for the 4 waveguides system, which are able to improve up to 76% and 42%, the execution time of the benchmarks. For the 16 waveguides system, the bests solutions are between the range from 8 (1,8) to 32 (2,16) channels. These allocations are able to enhance the execution time of the benchmark up to 77% and 74%, respectively. Fig. 7 shows the results for the normalized execution time of the SPLASH-2 benchmarks for different communication granularities.

TABLE II. COST OF THE CHANNEL ALLOCATOR

Network complexity	4 waveguides		16 waveguides	
Channel size (wg,wl)	Area (um2)	Power (mW)	Area (um2)	Power (mW)
(16,16)	NA[1]	NA[1]	6292	2.09
(8,16)	NA[1]	NA[1]	6629	2.21
(4,16)	6292	2.09	7304	2.45
(2,16)	6629	2.21	8781	2.95
(1,16)	7304	2.45	13538	4.45
(1,8)	8781	2.95	17022	5.06
(1,4)	13538	4.45	21385	5.33
(1,2)	17022	5.06	30110	5.86
(1,1)	21385	5.33	47560	6.92

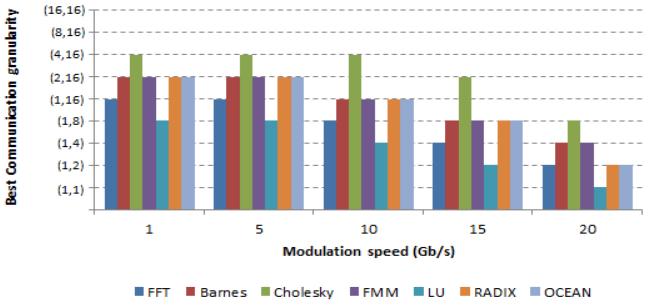

Fig. 8. Best channel size for different modulation speeds.

The values were normalized with the worst execution time for each application. The results show that for each application a different best channel bandwidth is found. The best communication granularities are in the interval from (1,2) to (2,16). The (1,2) channel size was the best for LU and (1,4) for FFT. For Barnes, FFM, RADIX and OCEAN the best channel size is (1,8). Finally for Cholesky, the best is (2,16). Moreover, the results show that applications characterized by huge percentage of writes and reads transactions (OCEAN, RADIX and BARNES) present execution curves that behaves as the hot spot traffic, where the execution time of (1,1) and (16,16) is similar.

D. *Modulation speed and communication channel size*

Modulation speed also affects the tuning of the channel size. Fig. 8 shows the variation of the best granularities according to rate at which data is injected to the optical network. Higher modulation speeds might result in smaller communication granularities. For data intensive applications, such as BARNES, FFM, RADIX and OCEAN, the channel size decreases from (2,16) channels at 1 Gb/s, up to (1,4) channels for BARNES and FFM and up to (1,2) channels for RADIX and OCEAN at 20Gb/s. Note that the decrease rate of the channel size depends on the optical MPSoC traffic. For example, the best channel size for Cholesky only starts to decrease after 10Gb/s, while for LU, it starts after 5Gb/s.

VI. Conclusions

From the system point of view, interface configuration is one of the most critical problems in optical interconnected Multi-processors Systems-on-Chip. The selection of the channel size (number of waveguides and wavelengths) dedicated to each request at design-time is advantageous because it generates simple channel allocator modules and small allocation times. Thus the design space exploration of the interface configuration is mandatory. In this work we provide for the first time an exploration of the communication channel size for an optical interconnected MPSoC. We show the huge impact of the communication channel size on the system performance, area and power. We also explore the effect on the modulation speed over the selection of the best channel size. We employ several synthetic and real application benchmarks. As future work we plan to implement strategies that guarantee the system performance while keeping the flexibility.

References

[1] A. Sheibanyrad, F. Pétrot, A. Jantsch. 3D integration for NoC-based SoC Architectures. Integrated Circuits and Systems. Ed. Springer. 2011, pp. 27-44.

[2] C. Battern, A. Joshi, V. Stojanovc. Designing Chip-Level Nanophotonic Interconnection Networks. Integrated Optical Interconnect Architectures for Embedded Systems. Embedded Systems. Ed. Springer. 2013, pp. 81-137.

[3] S. Le Beux, J. Trajkovic, I. O'Connor, G. Nicolescu, G. Bois, P. Paulin. System Level Exploration for the Integration of Optical Networks on Chip in 3D MPSoC Architectures. Integrated Optical Interconnect Architectures for Embedded Systems. Embedded Systems. Ed. Springer. 2013, pp. 241-262.

[4] S. Le Beux, H. Li, I. O'Connor, K. Cheshmi, X. Liu, J. Trajkovic, G. Nicolescu. Chameleon: Channel efficient optical network-on-chip. Proceedings of the conference on Design, Automation & Test in Europe, March 24-28, 2014, Dresden, Germany

[5] X. Wu, J. Xu, Y. Ye, Z. Wang, M. Nikdast, X. Wang. SUOR: Sectioned Undirectional Optical Ring for Chip Multiprocessor. ACM Journal on Emerging Technologies in Computing Systems, 2014.

[6] D. Vantrease, R. Schreiber, M. Monchiero, M. McLaren, N. Jouppi, M. Fiorentino, A. Davis, N. Binkert, R. Beausoleil, J. Ho. Corona: System Implications of Emerging Nanophotonic Technology. Proceedings of the 35th Annual International Symposium on Computer Architecture, p.153-164, June 21-25, 2008.

[7] R. Morris, A. Kodi, A. Louri. Dynamic Reconfiguration of 3D Photonic Networks-on-Chp for Maximizing Performance and Improving Fault Tolerance. IEEE/ACM 45th Anual 'nannualtional Symposium on Microarchitecture, 2012.

[8] J. Psota, et al. ATAC: Improving Performance and Programmability With on-Chip Optical Networks. Proceedings of IEEE International Symposium on Circuits and Systems, ISCAS, 2010.

[9] 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. In Design, AutomationTest in Europe Conference Exhibition, 2013.

[10] Z. Pan . Design for Manufacturability and Reliability for TSV-based 3D ICs. In 17th Asia and South Pacific Design Automation Conference (ASP-DAC), 2012, pp. 750-755.

[11] S. Koohi, M. Abdollahi, S. Hessabi. All-optical wavelength-routed NoC based on a novel hierarchical topology. In Proceedings 5th 2011 IEEE/ACM Int. Symp. Networks Chip, pp.97 -104.

[12] A. Krishnamoorthy, R. Ho, X. Zheng, H. Schwetman, J. Lexau, P. Koka, G. Li, I. Shubin, J. Cunningham. Computer Systems based on silicon photonic interconnects. In Proc. IEEE 97, 1337–1361, 2009.

[13] S. Le Beux, J. Trajkovic, I. O'Connor, G. Nicolescu, G. Bois, P. Paulin. Optical ring network-on-chip (ORNoC): Architecture and Design Methodology. In Proceedings of the conference on Design, Automation & Test in Europe, 2011, pp. 788-793.

[14] N. Binkert, B. Beckmann, G. Black, S. K. Reinhardt, A. Saidi; A. Basu, J. Hestness; D. R. Hower, T. Krishna, S. Sardashti, R. Sen, K. Sewell, M. Shoaib, N. Vaish, M. D. Hill, D. A. Wood. The gem5 simulator. SIGARCH Comput. Archit. News , 2011.

[15] N. Agarwal, T. Krishna, L. Peh, N. Jha. GARNET: A Detailed On-chip Network Model inside a Full-system Simulator. Proceedings of International Symposium on Performance Analysis of Systems and Software. 2009.

[16] S. C. Woo, M. Ohara, E. Torrie, J. P. Singh, A. Gupta. The SPLASH-2 Programs: Characterization and Methodological Considerations. In Proceedings of the 22nd International Symposium on Computer Architecture, pages 24–36, June 1995.

[17] J. Luo, C. Killian, S. Le Beux, D. Chillet, H. Li, I. O'Connor, O. Sentieys. Channel allocation protocol for reconfigurable Optical Network-on-Chip. In Proceedings of the SiPhotonics, Exploiting Silicon Photonics for energy-efficient heterogeneous parallel architectures. Hipeac, 2015.

[18] M. Ortín, L. Ramini, V. Viñals, D. Bertozzi. Capturing the Sensitivity of the Optical Network Quality Metrics to ts Network Interface Parameters. In Concurrency and Computation: Practice and Experience, 2014.