Scalable CMOS-compatible photonic routing topologies for versatile networks on chip

Alberto Scandurra
On Chip Communication Systems
STMicroelectronics
Catania, Italy
alberto.scandurra@st.com

Ian O'Connor
Lyon Institute of Nanotechnology
Ecole Centrale de Lyon
Lyon, France
Ian.Oconnor@ec-lyon.fr

***Abstract*— Optical network on chip (ONoC) architectures are emerging as potential contenders to solve both physical (routing, wire congestion) and performance (bandwidth, latency) issues in future computing architectures. In this work, we present a scalable and fully connected ONoC topology for multiple-core and heterogeneous SoCs. We show that it is possible, through careful design of network interfaces, to use the ONoC directly with existing protocols, while still exploiting specific optical properties and improving overall performance metrics, most notably that of congestion.**

I. INTRODUCTION

The shift to very high performance distributed Multi-Processor Systems-on-Chip (MPSoC) as mainstream computing devices is the recognized route to address, in particular, power issues by reducing individual processor frequency while retaining the same overall computing power. This rationale answers the need for flexible and scalable computing platforms capable of (i) achieving future required application performance in terms of resolution (audio, video and computing) and CPU power / total MIPS (real-time encoding-decoding, data encryption-decryption), and (ii) of working with multiple standards and with constrained power, which are both particularly important for mobile applications.

However, the move to such architectures requires organized high-speed communication between processors and therefore has an impact on the interconnect structure. It clearly relies upon the existence of an extremely fast and flexible interconnect architecture, to such a point that the management of communication between processors will become key to successful development. Aggregated on-chip data transfer rates in MPSoC, such as the IBM Cell processor [1], is critical and is expected to reach over 100Tb/s in the coming decade. As such, interconnects will play a significant role for MPSoC design in order to support these high data rates.

At the architectural level, networks on chip (NoC) overcome the limitations of bus-based platforms by providing each IP block, interfaced towards the network, with one or more reconfigurable channels of high-speed communication. NoC architectures are based on multiple data links interconnected by routers implementing packet switching for resource multiplexing. At the physical communication level, it is increasingly recognized that electrical interconnect will be highly inefficient in NoCs due to increasing power and silicon real estate concerns. One of the main replacement technologies currently under development consists of using integrated optical interconnect. Besides a huge data rate, optical interconnects also allow for additional flexibility through the use of wavelength division multiplexing. Exploring this aspect is necessary since it is not clear that a *direct* (single-wavelength) replacement of electrical links between switchboxes in a NoC topology by optical interconnect will achieve a significant performance gain, since this would require conversion between optical and electrical domains *at each switchbox*. Instead, through a shift in the routing paradigm (where the address of the target is not contained in the data packet but rather in the wavelength of the optical signal), it is possible to exploit this additional flexibility to design more intelligent interconnect systems, such as passive, wavelength-reconfigurable optical networks on chip (ONoC).

In section II we introduce the limit of classical electrical interconnect, and the need for an alternative solution. In section III an overview of a current NoC solution, the one developed by STMicroelectronics, is presented. Section IV details the architecture and principle of operation of the generic optical network on chip structure. Finally in section V, we cover the main communication scenarios for ONoCs.

It's important to point out that the focus of this paper is mainly on the topology of optical NoC, relying on the assumption that technology and design techniques allow to have an effective implementation of the physical layer, i.e. emitters, detectors and transport.

II. LIMIT OF ELECTRICAL INTERCONNECT

As design rules drop below 90 nm, a variety of challenges emerge such as RC delay, electromigration resistance, and heat dissipation exacerbated by increased chip power. The use of copper and thin barrier layers solves resistivity and electromigration problems but not for long due to electron scattering issues' increasing the apparent resistivity. Moreover, reliability issue with respect to an efficient diffusion barrier is a concern. Low k dielectrics allowing capacitance reduction have low thermal conductivity and hence poor heat dissipation

[1] This work was partially funded by FP7-ICT-2007-1-216405 "WADIMOS"

capability. Integration of copper and low k dielectrics is intensively studied worldwide [2].

Optical interconnects seem to be an alternative solution to overcome the issue of speed and power, providing much greater bandwidth, lower power consumption, decreased interconnect delays, resistance to electromagnetic interference and reduced signal crosstalk.

Photonic materials where light can be generated, guided, modulated, amplified and detected need to be integrated with standard CMOS integrated circuits in order to mix the information processing capability of electronics with the information transmission capability of photonics, providing a significant performance breakthrough within a cost effective engineering.

III. NETWORK ON CHIP SOLUTION OVERVIEW

The current ST NoC solution is based on a Network on Chip architecture called VSTNoC (Versatile STNoC), and evolves from the STBus approach [3]. It is in fact an interconnect system which has the same structure and functionality as the STBus, but uses a NoC-based protocol with appropriate interfaces and links. This approach enables higher performances and dramatic reductions in the number of pins and wires of the interconnect system, giving benefits in terms of area and facilitating rapid prototyping with FPGAs.

An STBus interconnect is composed of a set of building-blocks (nodes, converters and buffers) that can be cleanly assembled together in order to build almost any kind of architecture, from the simplest to the most complex one.

Figure 1. shows an interconnect built up with the VSTNoC, where network interfaces, nodes and buffers are used. In this figure we can see the uniformity in terms of both protocol (type) and bus size of the network with respect to the STBus; in fact all the required conversions are performed by the network interfaces where required, in order to adapt protocol, bus size and operation frequency to those of the network.

The VSTNoC solution belongs to the topology-dependent family. This means that, depending on the system topology (i.e. the number and type of initiators and targets of the system), the network topology can have different structures.

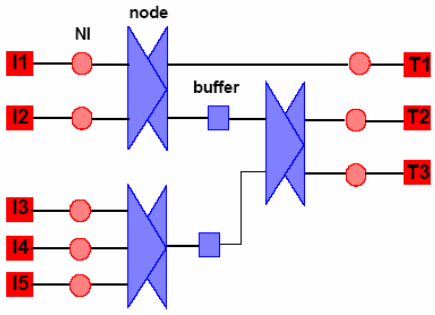

Figure 1. VSTNoC interconnect example

A. *VSTNoC Protocol*

The main features of the VSTNoC protocol are:

- a parametric header structure, the first field (IP_prot) of which identifies the protocol of the IP generating the traffic. According to the value of this field, the subsequent fields can differ in both meaning and size, depending on the IP native protocol;
- a NoC interface signal (aux/r_aux) carrying information about boundaries between possible elements characterizing different possible hierarchy levels of the IP native protocol (i.e. packet, chunk and message in STBus context, packet and burst in AMBA context);
- a flit identifier (flit_id) carrying information about the start and the end of a NoC transaction (a NoC transaction is a collection of NoC packets), determining the transaction or arbitration granularity (AG);
- an optional field in the response path carrying information about transaction status (r_flit_status), indicating whether errors have occurred, and which flits are affected.

B. *VSTNoC Transactions*

The VSTNoC transaction consists of the transmission of information from a traffic source to a destination according to a format that closely follows the one of usual network packets. It is the highest level transmission entity, marked by a start and an end, and can be chosen equivalent to an STBus message, chunk or even a single packet. The VSTNoC transaction is an atomic element, i.e. *it is not interruptible*.

A VSTNoC transaction is composed of VSTNoC packets, consisting of a header and a payload, the presence of which depends on specific conditions.

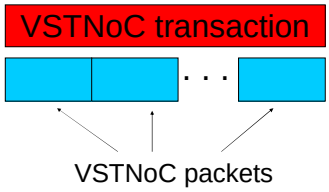

Figure 2. VSTNoC transaction structure

From a physical point of view, packets are split into basic units called flits (FLow control unITs). These represent the data link layer elements transmitted within a clock cycle in the case of synchronous transmission, or as asynchronous entities, the size of which is generally greater than the phyt (physical layer element) size. However it has been chosen equal to the phyt size in the VSTNoC context (i.e. in this work, flits and phyts are equivalent).

The flit is chosen to be sufficiently wide so as to be able to contain both data and byteenables transmitted over one single STBus cell in the request path. The following tables show the possible flit sizes that can be specified in an STNoC system for both the request and the response path, together with the different fields within a payload flit.

TABLE I. VSTNOC INTERCONNECT EXAMPLE

Flit size (bits)	Request flit field		Response flit field	
	be	data	be not used	r_data
36	<35:32>	<31:0>	-	<31:0>
72	<71:64>	<63:0>	-	<63:0>

The response flits are smaller than the request flits since in the response path there is no need to transport the byteenables signal, so fewer wires are required for the response interfaces.

C. VSTNoC Building Blocks

The VSTNoC communication system is based on the following building-blocks:

- Initiator Network Interface, responsible for IP to NoC traffic conversion and write posting response generation
- Target Network Interface, responsible for NoC to IP traffic conversion and internal errors (security and power down) management
- Node, responsible for buffering, arbitration, routing and wrong address errors
- Programming Module, allowing STNoC registers configuration
- Generic Converter, allowing to connect different NoC domains (with different flit size and/or frequency) and/or breaking long paths

In the next section, we will cover the description of the Optical Network On-Chip (ONoC), based on the VSTNoC.

IV. SCALABLE ONOC ARCHITECTURE

In an ONoC communication system, information is transmitted in the form of light, in opposition to the situation in classical electrical NoCs where the information is transmitted in the form of electrical charge (voltage levels on capacitors and currents for switching between voltage levels).

Communication relies on the ISO-OSI protocol stack, and can be seen as very close to the VSTNoC architecture, where the physical layer is replaced with a completely new one, exploiting optoelectronics in order to transmit information in form of light. The aim of this work is to demonstrate effective compatibility of the ONoC at the physical layer with the VSTNoC protocol.

The ONoC architecture consists of five main sets of building-blocks, as shown in Figure 3:

- **Initiator Network Interface** (**INI**): responsible for the conversion of the traffic generated by an initiator into a form suitable to be transmitted in form of light over the ONoC;
- **Transmitter**: responsible for the actual conversion of information from the electrical form into optical form, by means of information encoding for minimizing the power consumption by keeping the light emitter turned off as much as possible, serialization, emitter selection, emitter driving;
- **λ-Router (scalable passive integrated photonic routing structure)**: responsible for the actual propagation of optical information streams from sources to destinations;
- **Receiver**: responsible for the conversion of information from the optical form into electrical form, by means of photocurrent to voltage conversion, level adjustment, de-serialization, information decoding (for power consumption issue) and arbitration in case of multiple access from different traffic sources;
- **Target Network Interface** (**TNI**): responsible for the conversion of the traffic generated by the ONoC receiver into a form suitable to be received by the target.

Notice that INI and TNI are modules belonging to the electrical domain, while transmitter, receiver and l-router belong to the optical domain. Because of the serialization, the flit size in the electrical domain does not affect the optical network, but just the required storage at buffers in the electrical domain.

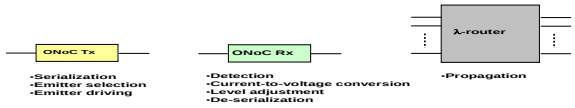

Figure 3. ONoC building blocks in optical domain

Such building-blocks can be assembled together to build proper on-chip communication architectures.

A. Principle of operation

An N×N ONoC, from a functional point of view, has the same behavior as an electrical N-port NoC: each initiator port (among N) can communicate simultaneously with one (or more, and possibly any number up to N) of N target ports. In this work, the quantity N represents the number of IP blocks to be connected through the communication structure; hence each IP block sends data through an initiator port and receives data through a target port. As previously mentioned, the ONoC is composed of a set of N transmitters and N receivers (one for each initiator port and target port respectively), and a scalable passive integrated photonic routing structure (λ-router). In this section, we will cover the principle of operation of this architecture and present results of physical and architectural evaluations from previous work.

Figure 4. shows an example of an 8×8 ONoC architecture.

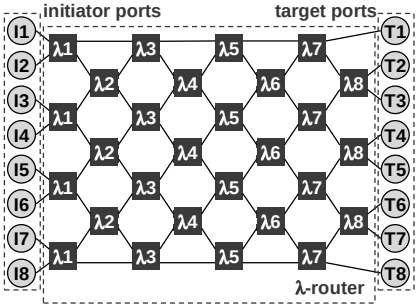

Figure 4. Full 8x8 ONoC topology schematic

In this representation, each initiator port I_i ($\forall i \in \{1,2,\ldots,8\}$) consists of a network interface (NI) and transmitter; and each target port T_j ($\forall j \in \{1,2,\ldots,8\}$) consists of a receiver and NI. Data is sent through the passive λ-router optically from each initiator to one or more targets by selecting a specific wavelength (for each initiator-target pair); in fact, only one physical path associated with a single wavelength exists between I_i and T_j. At any one time, a maximum of 8 (N) connections can exist in the network if each transmitter is equipped with a single, tunable-wavelength source; and a maximum of 64 (N^2) connections can exist in the network if each transmitter is equipped with N single- wavelength sources.

In the figure, each box containing λ_x represents a passive photonic component called an "add-drop filter" which can realize the key functionality of selecting and redirecting a signal based on its wavelength. There are many ways of realizing a photonic add-drop filter. In our work, we consider the use of passive microdisk resonators as shown in Figure 5. [3], for which the overall footprint can be considered to be approximately $10\times10\mu m^2$. Resonance in the individual microdisks occurs whenever the wavelengths of the optical signal carried by the neighboring waveguide corresponds to an integer number of lobes around the circumference of the microdisk, i.e. when the energy is distributed in the disk in *whispering gallery* modes. Because of this, the resonant wavelengths of a microdisk depend, for a given technology (and material parameters), on the radius of the microdisk.

As shown in the figure, the switching direction depends on the input wavelength λ and its relation to the resonant wavelength of the add-drop filter:

- when $\lambda=\lambda_n$ (within a given tolerance range depending on the quality factor of the microdisk) the signal will couple into the microdisk and then couple out into the waveguide in the same plane as the input. This is the *straight*, or *bar*, state.
- when $\lambda\neq\lambda_n$ the signal will propagate along the same waveguide and outputs in a different plane to the input according to the geometry of the waveguide. This is the *diagonal*, or *cross*, state.

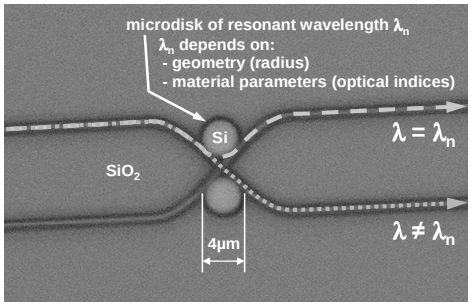

Figure 5. Si/SiO_2 microdisk-resonator based add-drop filter

When the WDM[2] technique is used, i.e. when multiple signals of various wavelengths are injected at the input (which is usually the case to increase the global throughput of the network), a cumulative state occurs, where individual signals simultaneously obey the routing characteristics of the add-drop filter according to their individual wavelengths. Because of this property and the fact that the four add-drop ports can be used simultaneously, a contention-free network can be built.

The overall passive λ-router network consists of N stages of alternately N/2 and (N/2)–1 add-drop filters (or, more generically, routing elements). Using microdisk resonators, the overall area required for the 8×8 passive network is around $3000\mu m^2$. The path followed by the optical signal in the overall network shown in Figure 4. depends only on the wavelength and can be obtained by equation (1). As an example, if the block at initiator port I_3 is to communicate with the block at target port T_5, then I_3 must send data through the λ-router with wavelength λ_1. It is thus clear that each IP block can "reconfigure" its communication paths by using different wavelengths.

The matrix shown in equation (1) displays two interesting properties. Firstly, it is symmetrical around both diagonals. This means that the set of communication properties of the top half of the network is the flipped mirror image of that of the bottom half of the network; and that the return path for communication is exactly the same as the transmission path. The second noteworthy property is the existence of non-resonant wavelengths in certain communication paths (shown in bold in the matrix). While specific wavelengths have been assigned in the matrix to these communication paths, *any* wavelength (other than the wavelengths used by the other communication paths) can be used. This is the case since these communication paths do not actually pass through a routing element corresponding to the assigned wavelength at all – they cross the (N/2)-1 routing stages at the top or at the bottom of the network and thus only pass through a waveguide, rather than a resonant routing element. In the full ONoC, the unused wavelengths are assigned to these communication paths in order to exploit the resources – however this property can also be exploited to reduce the number of wavelengths used [5] [6].

$$\begin{bmatrix} T_1 \\ T_2 \\ T_3 \\ T_4 \\ T_5 \\ T_6 \\ T_7 \\ T_8 \end{bmatrix} = \begin{bmatrix} \lambda_4 & \lambda_5 & \lambda_3 & \lambda_6 & \lambda_2 & \lambda_7 & \lambda_1 & \boldsymbol{\lambda_8} \\ \lambda_5 & \lambda_6 & \lambda_4 & \lambda_7 & \lambda_3 & \lambda_8 & \boldsymbol{\lambda_2} & \lambda_1 \\ \lambda_3 & \lambda_4 & \lambda_2 & \lambda_5 & \lambda_1 & \boldsymbol{\lambda_6} & \lambda_8 & \lambda_7 \\ \lambda_6 & \lambda_7 & \lambda_5 & \lambda_8 & \boldsymbol{\lambda_4} & \lambda_1 & \lambda_3 & \lambda_2 \\ \lambda_2 & \lambda_3 & \lambda_1 & \boldsymbol{\lambda_4} & \lambda_8 & \lambda_5 & \lambda_7 & \lambda_6 \\ \lambda_7 & \lambda_8 & \boldsymbol{\lambda_6} & \lambda_1 & \lambda_5 & \lambda_2 & \lambda_4 & \lambda_3 \\ \lambda_1 & \boldsymbol{\lambda_2} & \lambda_8 & \lambda_3 & \lambda_7 & \lambda_4 & \lambda_6 & \lambda_5 \\ \boldsymbol{\lambda_8} & \lambda_1 & \lambda_7 & \lambda_2 & \lambda_6 & \lambda_3 & \lambda_5 & \lambda_4 \end{bmatrix} \begin{bmatrix} I_1 \\ I_2 \\ I_3 \\ I_4 \\ I_5 \\ I_6 \\ I_7 \\ I_8 \end{bmatrix} \quad (1)$$

B. *Evaluated performance metrics*

In prior work [6], a 4×4 passive λ-router was fabricated and measurements show that its operation agrees with the theory. Resonant wavelengths were measured between 1547-1583nm for Si/SiO_2 microdisks of radii from 1.0-2.5μm. The minimum

[2] Wavelength Division Multiplexing

free spectral range (FSR[3]) was measured to be 50nm, and quality factors around 500-800.

In parallel work, the design of a 16×16 ONoC virtual prototype was carried out at various abstraction levels using a top-down approach [8] from architecture to physical design, enabling an accurate estimation of various performance metrics. The source and detector characteristics were extracted from III-V device data, and transistor-level interface circuits sized with a 0.13 μm CMOS technology. In this context, the ONoC can achieve a data rate of up to 3.2Gb/s per port with a latency of 420ps and power consumption of 10mW per unidirectional link. The ONoC data rate is in fact limited by the interface circuits, mainly at the receiver. The SERDES circuits contribute greatly to power consumption at these frequencies.

More recently in [9], the impact of the low latency and absence of contention in the ONoC interconnect architecture was assessed for an 8-processor SoC running an MPEG-4 algorithm. When comparing a 100MHz ONoC against 200MHz STBus [10] and 2- and 5-CCL[4] crossbars, the ONoC demonstrated speedup factors of between 1.5 and 3.2, i.e. better performance, in terms of processing time, than any traditional electrical interconnect, even at half the operational frequency.

V. ONoC ARCHITECTURE COMMUNICATION SCENARIOS

In this section, we cover the uses of ONoC in actual communication scenarios.

A. Communication scenarios

The optical waveguides within the ONoC are bidirectional. However, two-way communication between 2N IP blocks over a single ONoC is not feasible since this would require optical detectors and sources with identical wavelength selectivity to lie on the same waveguide with no interaction – this is clearly impossible. Additionally in this configuration there can be no communication among IP blocks which have been assigned ports situated (physically) on the same side of the passive routing network. In fact there are two scenarios for the use of the N×N ONoC, both using the ONoC for communication in a single direction only:

- in the first scenario, shown in Figure 6. (a) for 8 IP blocks, we consider that each IP block is assigned a pair of initiator/target ports. This leads to total connectivity between all N IP blocks, and to the non-use of wavelengths corresponding to communication paths I_i-T_j when i=j.

- in the second scenario, shown in Figure 6. (b) for 8 IP blocks, we consider that two identical (N/2)×(N/2) ONoCs are used for request/response type communications between two sets of N/2 IP blocks. In this case, no communication is possible between IP blocks in the same set, but this scenario does lead to reduced requirements on the overall number of wavelengths and routing elements.

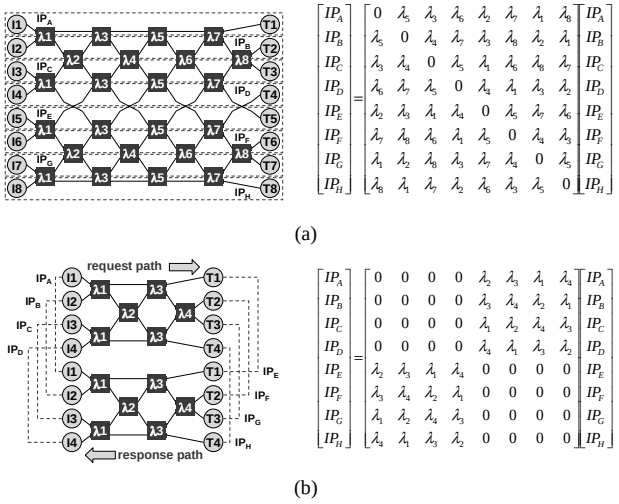

Figure 6. Communication scenarios and corresponding connectivity matrices for ONoC in 8-IP block scenarios (a) single 8×8 ONoC for total connectivity between 8 IP blocks (b) 2 4×4 ONoCs for request/response connectivity between 2 groups of 4 IP blocks

In Table II, a comparison is made between various performance metrics for each scenario. These represent extremes for (a) total connectivity and (b) balanced communication between groups of IP blocks of equal numbers. In practice, it is unlikely that the required system connectivity will fall into either of these scenarios. However, the total connectivity scenario represents the default or reference scenario, while the grouped connectivity scenario makes clear that if total connectivity is not required in the system, significant reductions in complexity can be achieved.

TABLE II. COMPARISON BETWEEN PERFORMANCE METRICS FOR TOTAL CONNECTIVITY AND GROUPED CONNECTIVITY ONOC SCENARIOS

	(a) Total connectivity	**(b) Grouped connectivity**
IP blocks	N	N
Connections	N(N-1)	$(N/2)^2$
Required wavelengths per IP block n_λ	N-1	N/2
Number of routing elements n_r	N(N-1)/2	N(N-1)/4

B. Physical considerations

The comparisons mentioned in Table I are important for several reasons. Firstly, the number of routing elements n_r impacts directly on the overall size and complexity of the passive routing network. The size of the photonic communication layer is limited by the size of the CMOS chip. If several parallel λ-routers can fit into this area, then data rate could be increased (or power consumption reduced by running at a lower clock frequency).

Secondly, the required number of wavelengths n_λ per IP block will impact directly on the number of transmitters (and

[3] FSR is defined as the difference between resonant wavelengths of a passive resonator. In the Si/SiO_2 microdisk resonators, FSR ≈ 50nm.

[4] Clock Cycle Latency

sources and wavelength multiplexers) and receivers (and wavelength demultiplexers and detectors) per IP block. The schematic of the transmitter structure and corresponding geometrical representation for the set of microdisk laser sources is shown in Figure 7. (a) and Figure 7. (b) respectively.

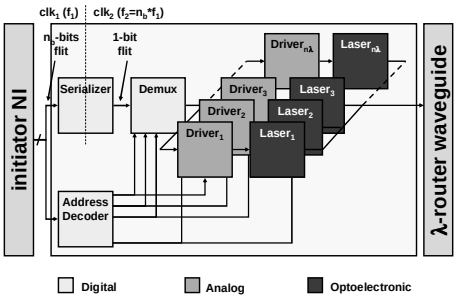

(a)

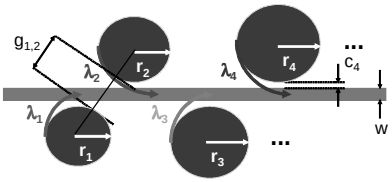

(b)

Figure 7. n_λ-laser source transmitter structure (a) schematic (b) corresponding geometrical representation for the set of microdisk laser sources

Since the laser source drivers are based on current modulation schemes, each source costs, in terms of static and dynamic power consumption, its bias current and modulation current respectively. As a consequence, the overall static and dynamic power consumption increases linearly with n_λ. In terms of the geometry and its impact on the size of the transmitter on the photonic layer, its area A_t can be expressed as

$$A_t = \left((n_\lambda - 1)\sqrt{(2\bar{r} + \bar{g})^2 - (w + 2(\bar{c} + \bar{r}))^2} + 2\bar{r} \right)(2(2\bar{r} + \bar{c}) + w) \quad (2)$$

where

$$\bar{c} = \frac{\sum_{n=1}^{n_\lambda} c_n}{n_\lambda} \; ; \; \bar{r} = \frac{\sum_{n=1}^{n_\lambda} r_n}{n_\lambda} \; ; \; \bar{g} = \frac{\sum_{n=1}^{n_\lambda} g_{n,n+1}}{n_\lambda}$$

and c_n represents the nominal source$_n$-waveguide distance (between 0.4-0.6µm), r_n the nominal microdisk laser radius (between 1-10µm), $g_{n,n+1}$ the minimum source-source spacing (typically 3µm), and w the waveguide width (under 1µm).

At the target end, each IP block requires a separate receiver path for each wavelength received, in order to identify the origin of each incoming data flit and also in order to be able to buffer flits incoming simultaneously from different initiator ports. The schematic of the receiver structure and corresponding geometrical representation for the set of microdisk demultiplexers and broadband photodetectors is shown in Figure 8. (a) and Figure 8. (b) respectively.

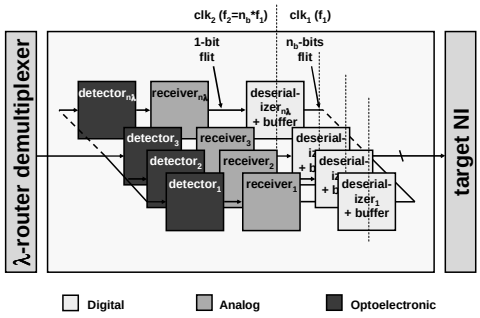

(a)

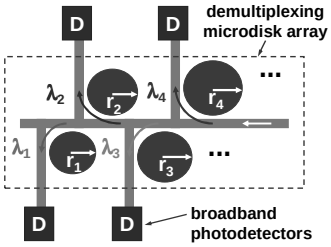

(b)

Figure 8. n_λ-demultiplexing receiver structure (a) schematic (b) corresponding geometrical representation for the set of microdisk demultiplexing elements and broadband photodetectors

Finally, as shown in Figure 9. , since the maximum WDM window is approximately equivalent to the FSR of the microdisk resonators, a larger number of wavelengths will also lead to more stringent constraints on the selectivity (Q factor) of each resonator, and on the accuracy of the lithography techniques used to define the radius (and resonant wavelength) of each passive microdisk in the λ-router. With current process technology characteristics, a maximum of around 16 distinguishable and stable wavelengths can be achieved.

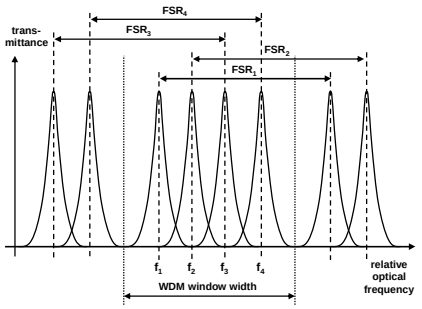

Figure 9. Relationship between microdisk resonator free spectral range and WDM window width

VI. CONCLUSIONS

In this work, we have described a scalable and fully connected N×N ONoC topology compatible with existing NoC paradigms. We have covered the main metrics that can be extracted from various communication scenarios for fully connected topologies, and in particular constraints on source

wavelength accuracy and passive filter selectivity depending on the number of required wavelengths, and power and area issues depending on the number of active and passive devices.

REFERENCES

[1] M. Gschwind, "Chip multiprocessing and the Cell broadband engine," *ACM Int. Conf. Computing Frontiers*, Ischia, Italy, May 2006

[2] L. Pavesi, G. Guillot (Eds.), "Optical Interconnects", Springer (2006)

[3] M. Coppola, C. Pistritto, R. Locatelli and A. Scandurra STNoC™: An Evolution Towards MPSoC Era", NoC Workshop, march2006, DATE 2006

[4] 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, no. 7, pp. 1447-9, July 2005

[5] I. O'Connor *et al.*, "Reduction methods for adapting optical network on chip topologies to specific routing applications," *Proc. Design of Circuits and Integrated Systems (DCIS)*, November 12-14 2008, Grenoble, France

[6] L. Zhang *et al.*, "Generalized Wavelength Routed Optical Micronetwork in Network-on-chip", Proc. 18th Int. Conf. Parallel and Distributed Computing and Systems (IASTED), pp. 698-703, Dallas, TX, Nov. 2006

[7] A. Kazmierczak *et al.*, "Design and characterisation of optical networks on chip," *Proc. Eur. Conf. Integrated Optics (ECIO)*, Grenoble, France, April 2005

[8] M. Briere *et al.*, "Heterogeneous modelling of an Optical Network-on-Chip with SystemC," *Proc. Rapid System Prototyping Workshop (RSP)*, June 2005

[9] M. Brière *et al.*, "System Level Assessment of an Optical NoC in an MPSoC Platform," *Proc. Design Automation and Test in Europe (DATE)*, Nice, France, April 2007

[10] A. Bona, V. Zaccaria, and R. Zafalon, "System level power modeling and simulation of high-end industrial network-on-chip," *Proc. Design, Automation and Test in Europe (DATE) Conference and Exhibition*, Munich, Germany, February 2004