



HAL
open science

Min/max time limits and energy penalty of communication scheduling in ring-based ONoC

Joel Ortiz Sosa, Cedric Killian, Hamza Ben Ammar, Daniel Chillet

► To cite this version:

Joel Ortiz Sosa, Cedric Killian, Hamza Ben Ammar, Daniel Chillet. Min/max time limits and energy penalty of communication scheduling in ring-based ONoC. NoCArc 2020, Oct 2020, On-line, France. hal-03032687v1

HAL Id: hal-03032687

<https://inria.hal.science/hal-03032687v1>

Submitted on 1 Dec 2020 (v1), last revised 1 Feb 2021 (v3)

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.

Min/max time limits and energy penalty of communication scheduling in ring-based ONoC

Joel ORTIZ SOSA
Univ Rennes, Inria, IRISA
France-22305 Lannion
joel.ortiz@inria.fr

Cedric KILLIAN
Univ Rennes, Inria, IRISA
France-22305 Lannion
cedric.killian@inria.fr

Hamza BEN AMMAR
La Rochelle University, L3i
France-17000 La Rochelle
hamza.ben_ammam@univ-lr.fr

Daniel CHILLET
Univ Rennes, Inria, IRISA
France-22305 Lannion
daniel.chillet@inria.fr

Abstract—Recent advances in the photonics devices integration bring ONoC as a bridge future for communication media in the MPSoC domain. As ONoC can support Wavelength Division Multiplexing (WDM) technique, communications between cores can be improved through allocation of one or several wavelengths for each communication. However, WDM introduces wavelength crosstalk, requiring to increase the laser power to provide accurate communication between cores. Thus, for the designer, exploring this design space (execution time vs power consumption) is not an easy task due to a large number of wavelength allocation combinations. The contribution presented in this paper proposes to evaluate the two extreme bounds of this design space considering the different communication scenario. To address this problem, we model the wavelength allocation by two different objective functions to compute the bounds in terms of execution times. Furthermore, from an accurate model of crosstalk between the wavelengths, we compute the energy penalty for each communication scenario. The results presented in this paper highlight the execution time and energy consumption tradeoff, and the opportunity for communication optimisation thanks to an efficient use of WDM technique.

I. INTRODUCTION

High-performance computing systems are generally based on multi-core architectures embedding hundreds of cores on a single chip. The design of such Multi-Processor System-on-Chip (MPSoC) is very challenging, not only for the computation part, but also for the communication media. Indeed, implementing applications on such architectures results in the different tasks being distributed onto the cores and adding more and more data transfers between these cores. To support these data transfers, a large number of classical Network-on-Chips (NoC) were designed and proposed in the literature [3]. However, these NoCs suffer from some limitations due to the electrical interconnect characteristics: capacitive and inductive coupling [8], interconnect noise and propagation delay. Thus, new on-chip interconnect solutions are highly desirable to overcome these limitations.

In parallel, recent advances in the integration of photonic on silicon allow the conception of optical NoCs [1] making them good candidates for on-chip communications. This technology provides high bandwidth for data transfers and can be a very interesting alternative to bypass the bottleneck induced by classical NoCs.

Among the different ONoC architectures, Multi-Writer Multi-Reader (MWMR) ONoCs offer the most scalable architectures as the communication medium is shared for the data transfers [13]. This scalability is increased with Wavelength Division Multiplexing (WDM) allowing multiple communications in parallel supporting by different wavelengths [6].

However, as for any shared communication medium, the scheduling of communications is a tough part and clearly impacts the application performance. Hence, the communication scheduling in ONoCs has been addressed to solve different purposes like minimizing the number of wavelengths used in the same portion of a waveguide [5], or proposing a trade-off between energy and execution time [11]. The main difficulty is to propose a method to extract the possible communication scenario to evaluate the trade-off between the execution time of an application and the energy consumption overhead needed to counter the crosstalk induced by using WDM technique.

This paper tackles this topic, providing an optimal solution. Besides, to the best our of knowledge, this study is the first to propose a formal model allowing to compute the optimal time extrema of wavelength scheduling in a ring-based WDM MWMR ONoC with an evaluation of crosstalk overhead.

The rest of the paper is organised as follow. Section II presents the communication scheduling impact on the execution time and the crosstalk power penalty. Section III presents the communication allocation and scheduling formalisation developed to find the different global communication solutions. The crosstalk model used for the evaluation of the crosstalk power penalty is also presented in this section. Section IV presents the results provided by the resolution of the problem. Section V describes the state of the art regarding existing methods to handle channel scheduling in ring based ONoCs. Finally, Section VI concludes this paper.

II. COMMUNICATION SCHEDULING

The key concepts investigated in this paper are illustrated in Figure 1 (the legend is presented in Fig.1.g). Applications are represented as a Directed Acyclic Graph (DAG) as depicted in Fig.1.a. In this example, we assume that tasks t_0 , t_1 , t_2 and t_3 are respectively mapped onto processors p_0 , p_1 , p_4 , and p_8 , respectively located in clusters c_0 , c_0 , c_1 and c_2 (see Fig.1.b). Communications $Com_{0 \rightarrow 2}$, $Com_{1 \rightarrow 2}$ and $Com_{2 \rightarrow 3}$

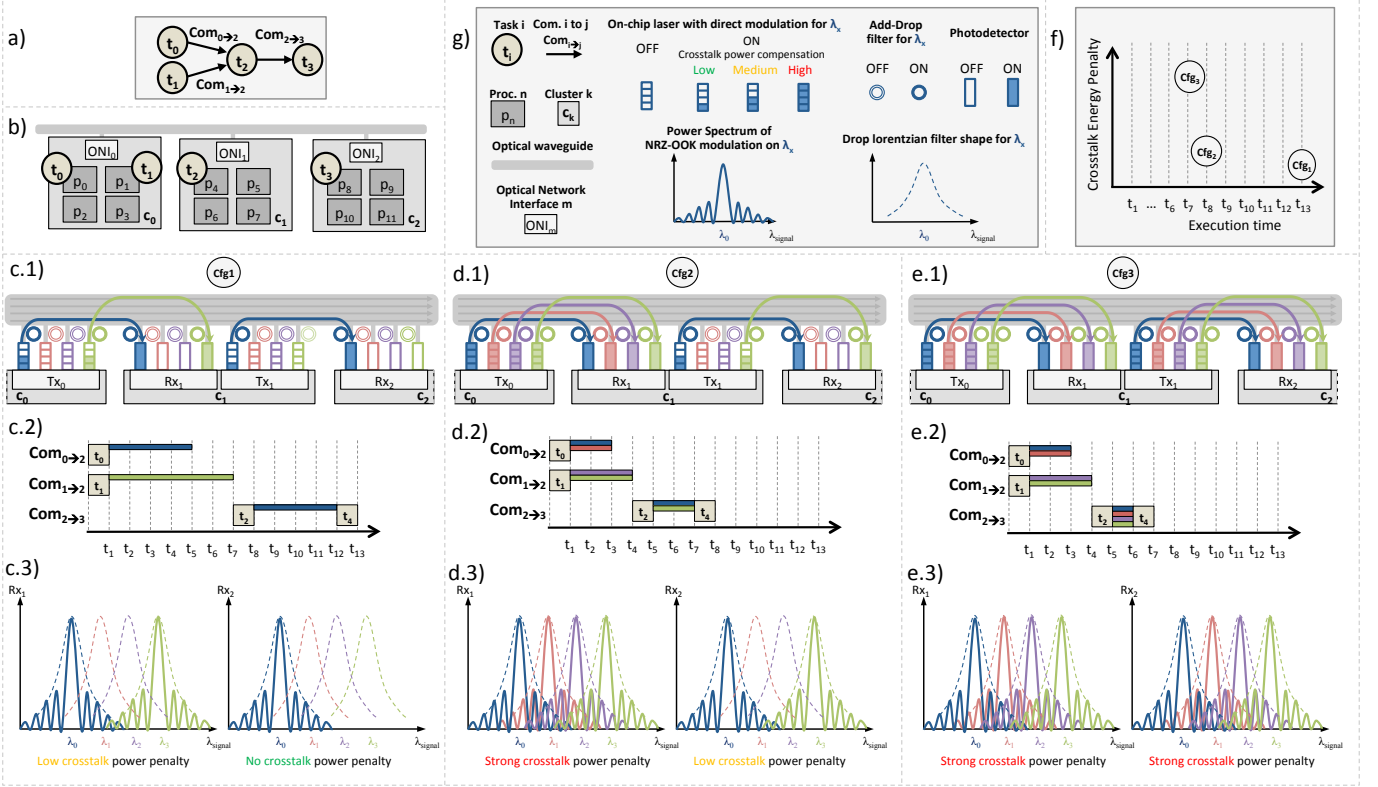


Fig. 1. Global concept of communication scheduling for Optical Network on Chip.

are implemented using an ONoC which can be configured according to execution performance and energy requirements. Each cluster is associated with an Optical Network Interface (ONI) containing a Transmitter (Tx) and a Receiver (Rx). The Txs integrate lasers to modulate data, and a Micro Ring (MR) resonator for each laser to inject the optical signal into the waveguide (if the MR is 'on'). The Rxs integrate MR to eject the associated optical signal from the waveguide when the MR is 'on'. Finally, photodetectors allow converting an optical signal into an electrical one. From the simple application given in Fig.1-a and if we consider a waveguide supporting 4 wavelengths, several communication configurations are possible and the following illustrates three of them:

- Low power configuration (Cfg1): in this configuration, illustrated in Fig.1.c.1, a single wavelength is allocated to each communication. Wavelengths λ_0 (blue) and λ_3 (green) in Tx_0 are allocated for $Com_{0 \rightarrow 2}$ and $Com_{1 \rightarrow 2}$, respectively, and λ_0 (blue) is also used in Tx_1 for $Com_{2 \rightarrow 3}$. We assume a direct OOK modulation, i.e. no modulator is needed. The signals propagate along the waveguide until they reach their destination: signals from Tx_0 at wavelength λ_0 and λ_3 are ejected in cluster c_1 , while λ_0 signal from Tx_1 is ejected at c_2 . Since λ_0 is used between c_0 to c_1 to support $Com_{0 \rightarrow 2}$, it is reused to implement $Com_{2 \rightarrow 3}$ between clusters c_1 and c_2 , thus maximizing the wavelength occupation in the waveguide. The Fig.1.c.2 illustrates the task and communication schedule. Processor p_4 in cluster c_1 starts executing t_2 once

all the data have been received from p_0 and p_1 and sends the processed data to p_8 in c_2 . The execution of the DAG in this configuration takes 13 clock cycles. In this example, we assume that the targeted BER is reached for all communications. The Fig.1.c.3 depicts the power spectrum of signals received at Rx_1 and Rx_2 . The NRZ-OOK modulated signals are represented as cardinal sines, while the filter MRs have Lorentzian filter shapes. For the communications between c_0 and c_1 , the two used wavelengths are the most distant, hence the crosstalk between λ_0 and λ_3 is limited, but the associated lasers must increase a bit the emitting power to compensate the signal degradation, named crosstalk power penalty (see Section III-C for more detail on crosstalk). Regarding the communication between c_1 and c_2 , as only one wavelength is received at Rx_2 , no crosstalk penalty has to be compensated. It is worth mention that the power of laser also needs to compensate for propagation losses (not represented here).

- Intermediate configuration (Cfg2): for this configuration more bandwidth is allocated for each communication. $Com_{0 \rightarrow 2}$, $Com_{1 \rightarrow 2}$ and $Com_{2 \rightarrow 3}$ are respectively supported by the wavelength pairs (λ_0, λ_1) , (λ_2, λ_3) and (λ_0, λ_3) . This allocation leads to a reduction of execution time in 8 clock cycles (see Fig.1.d.2). However, in Rx_1 , as all wavelengths are used, the crosstalk penalty is at maximum, hence the source lasers have to compensate a strong crosstalk power penalty. Regarding the Rx_2 , as the two used wavelengths

are the most distant, the crosstalk power penalty is low.

- Fastest configuration (c_{fg}) : for this configuration, the bandwidth utilization is at maximum. Compared to (c_{fb}) , $Com_{2 \rightarrow 3}$ is supported by the 4 available wavelengths. This allocation leads to the fastest execution of the DAG in 7 clock cycles but provides the hugest crosstalk (see Fig.1.e.2 and Fig.1.e.3) which needs to be compensated in each Tx .

The Fig.1.f plots the execution time versus the crosstalk energy penalty for each configuration. The previously presented configurations belong to a Pareto front. The configurations (c_{fg}) and (c_{fb}) are respectively the two bounds: the slowest and the fastest execution configurations. Based on the complexity of the application, the number of solutions on the Pareto front can be difficult to determine, as well as the front itself. As mentioned in the introduction, numerous methods have been developed to optimise the communication scheduling in ONoC. However, no mathematical formulation has been proposed considering MWMR and multiple wavelengths per communication. Hence, the state of the art solutions require a huge effort to explore all the solutions to determine the Pareto front. This paper tackles this problem and allows us to compute the two extrema solutions of this Pareto front. From these information, the designer can evaluate the optimisation opportunity of the wavelength allocation for a specific application using ONoC.

III. FORMALISATION OF COMMUNICATION SCHEDULING

This section presents the wavelength allocation problem, defined with two optimisation functions to extract the bounds in terms of execution time. From this information, and from the crosstalk model (presented in Section III-C), we compute the energy penalty of each bound, which finally gives us the Pareto front of the wavelength allocation design space.

A. Execution time bounds computation

As previously explained, the objective is to compute the two extrema solutions in terms of execution time. The first one, defined by the fastest execution time, is produced by an optimal wavelength allocation, while the second bound is produced by an allocation with only one wavelength for each communication. In this context, the problem addressed includes wavelength allocation and tasks scheduling. Before modelling the optimisation problem, let's define the following variables:

- $Tasks = \{T_i\}$ defines the set of tasks T_i to schedule. Each T_i is defined by the variables $\{Ts_i, Td_i, Te_i\}$ defined below.
- Nt is an integer defining the tasks number in the set $Tasks$.
- Td_i is an integer defining the execution time of the task T_i .
- $Clusters = \{P_j\}$ defines the set of clusters P_j included in the architecture. A cluster P_j is a set of processors which supports the execution of a subset of $Tasks$.
- $ONI = \{ONI_j\}$ defines the set of Optical Network Interface ONI_j included in the architecture. Each cluster P_j has its own ONI_j .
- Op is an integer defining the number of ONIs.

- $Com_{i,j}$ is a boolean variable defining the communication needs between the tasks. If $Com_{i,j} = 1$, it means that a communication exists from tasks T_i to T_j .
- $Data_{i,j}$ is an integer variable defining the volume of data transferred from the tasks T_i to T_j .
- M_i is an integer variable defining the task mapping. $M_i = j$ means that T_i is mapped on the cluster P_j .
- $SC_{k,Com_{i,j}}$ is a boolean variable equal to 1 if a communication $Com_{i,j}$ travels the waveguide segment S_k . A waveguide segment S_k is the part of waveguide between two ONIs. Typically, S_1 is the segment between ONI_1 and ONI_2 . Note that the static mapping of tasks (M_i) enables to pre-compute the values $SC_{k,Com_{i,j}}$.
- NW is an integer variable defining the maximum number of wavelengths supported by the waveguide.
- B is an integer variable defining the bandwidth for the transfers of data on one wavelength.
- t defines the execution cycles.

All the previous variables define the architecture, the mapping and characteristics of tasks. To complete the optimisation problem definition, decision variables must be defined, they are the following:

- $Alloc\lambda_{i,j}$ is the integer variable defining the number of wavelengths allocated for the communication between tasks T_i and T_j .
- Ts_i and Te_i are respectively the start time and the end time of the task T_i .

From these variables, we can define the two optimisation functions providing the execution time bounds. The first optimisation function enables to extract the solution with the minimum execution time for an application. This objective function is given by the following equation:

$$Min \quad (max\{Te_i\}) \quad \forall T_i \in Tasks \quad (1)$$

This expression minimises the end of application execution. Solving this problem provides a solution with the optimal number of wavelengths allocated in order to reduce the communication times.

The second objective function is necessary to find the maximum execution time, obtained when the number of wavelengths for each communication is equal to 1. This objective function is defined by:

$$Min \quad \sum_{i=0}^{Nt} \sum_j^{Nt} Alloc\lambda_{i,j} \quad (2)$$

For both optimisation functions, three constraints are defined.

- The first constraint, given as c_1 (Eq.3), is very simple and formalises the links between the start time, end time and the execution time of a task T_i .

$$c_1 : \quad Te_i = Ts_i + Td_i \quad (3)$$

- The second constraint, given as c_2 (Eq.4), verifies that the start time of task T_i is greater than the sum between i) the max of the end times of all the predecessor task T_j of T_i and ii) the time to transfer the data between the tasks T_j and T_i .

This time depends on the number of wavelengths allocated for the communication, more allocated wavelengths leads to smaller communication time.

$$c_2 : Ts_i = \max \left\{ Te_j * Com_{j,i} + \frac{Data_{j,i}}{Alloc\lambda_{j,i} * B} \right\}_{\forall Te_j} \quad (4)$$

Note that if there is no communication between tasks T_i and T_j , $Te_j * Com_{j,i} + \frac{Data_{j,i}}{Alloc\lambda_{j,i} * B}$ is equal to 0.

- The third constraint, given as c_3 (Eq.5), is the more complex part of the problem. This constraint verifies that, for a waveguide segment S_k and for a specific cycle t , the sum of wavelengths allocated to support the communications is not greater than the number of wavelengths available in the waveguide. For that, we compute, for each waveguide segment S_k , the sum of wavelengths used by each communication travelling this segment at the same time, and this sum must be smaller or equal than the maximum number of wavelengths.

$$c_3 : \forall \{k, t\} \mid k = 0, \dots, Op - 1 \\ \wedge t = 0, \dots, \sum_i \left(Te_i + \frac{\sum_j Data_{i,j}}{B} \right) \\ \sum_{i=0}^{Nt-1} \sum_{j=0}^{Nt-1} SC_{k,Com_{i,j}} * Com_{i,j} * (Te_i < t) * (Ts_j > t) \\ * Alloc\lambda_{i,j} \leq NW \quad (5)$$

To describe these two problems, we use the language OPL (Optimisation Programming Language) and Cplex solver to explore the solutions and to extract the bounds of execution times. The two objective functions concern discrete integer decision variables and the optimisation only concerns one objective. Then the problem can be classified as Integer Linear Program.

B. Computation of crosstalk penalty

As the mapping of a task is statically defined and from the scheduling produced by the above optimisation functions, we can compute the crosstalk penalties which appear at each receiver R_x . We consider each communication $Com_{i,j}$ and we compute the number of wavelengths used from transmitter to receiver ONIs of this communication. For that, we extract all the communications travelling at least on one segment of communication $Com_{i,j}$ and for which a time overlap exists. If $Com_{i,j}$ used n wavelengths to support the communication from tasks T_i to T_j , then each of the n wavelengths is crosstalked with the $n - 1$ other wavelengths, it defines *AutoCrosstalk*, AC , given by Eq.6. Furthermore, each of these n wavelengths are crosstalked by all the other wavelengths supporting the other communications using the same part of the waveguide and with a time overlapping, it defines *InterCrosstalk*, IC , given by Eq.7.

$$AC_{i,j} = Alloc\lambda_{i,j} * (Alloc\lambda_{i,j} - 1) \quad (6)$$

$$IC_{i,j} = Alloc\lambda_{i,j} * \sum_{k,l \mid (k,l) \neq (i,j)} (Alloc\lambda_{k,l} \\ * TimeOverlap(Com_{i,j}, Com_{k,l}) \\ * SegmentOverlap(Com_{i,j}, Com_{k,l})) \quad (7)$$

With *TimeOverlap*, respectively *SegmentOverlap*, functions returning 1 if the two communications share at least one communication cycle, respectively one segment of the waveguide. The functions return 0 elsewhere.

From these two values that count the number of wavelengths which crosstalk together, we compute the power penalty based on the crosstalk model defined in the next section.

C. Crosstalk model

As introduced in Section II, wavelength crosstalk appears when two wavelengths are close to each other and when an MR extracts power from its own resonant wavelength plus power from the neighbour wavelengths. The signals propagating within the waveguide have cardinal sine shape due to the NRZ-OOK modulation used. We use the model of [2] to estimate the crosstalk. In this model, we consider especially: the wavelength frequency of each channel, the frequency spacing between channels, FWMR of MR, and the bit-rate of transmission (please refer to [2] for more detail on the model). This model allows us to compute the crosstalk power penalty (XPP) for each signal. This value is the power compensation at the source lasers that is required to maintain the signal quality (Bit Error Rate) due to wavelength crosstalk.

Fig. 2 presents the crosstalk power penalty which impacts each wavelength in presence of all the other wavelengths. These power penalties are computed by considering a constant FSR, 0.8 nm in our case, and with an optimal repartition of wavelengths into the FSR interval. It means that when the waveguide includes only two wavelengths, the FWMR is at the maximum value (equal to FSR), and the crosstalk between the two wavelengths is near 0 dB , while the power penalty is near 4 dB for waveguide including 16 wavelengths and an FWMR equal to $FSR/15 = 0.053 \text{ nm}$. For the computation of crosstalk penalty, we consider the maximum value of the power penalty $XPP_{max_{Nw}}$ given in the Fig.2. For example, for a waveguide supporting 16 wavelengths (green curve of Fig.2), $XPP_{max_{16}} = 5.2 \text{ dB}$. Finally, the global crosstalk energy penalty ($GXEP$) of an application is then computed by Eq.8 which takes account of the communication duration (in cycles) which is the time interval when the laser power must be increased.

$$GXEP = \sum_{i,j} (AC_{i,j} + IC_{i,j}) * XPP_{max_{Nw}} * (Ts_j - Te_i) \quad (8)$$

As we don't include clock frequency in this equation, the unit of $GXEP$ is given in $\text{dB} * \text{cycles}$.

IV. RESULTS

We evaluate our contribution with a set of synthetic task graphs generated with our own task generator. This task generator is very similar to TGFF tool, but it includes communication volume between tasks. For this experimental setup, the parameters used to generate the task graphs are the following: number of tasks from 6 to 12 tasks, parallelism between tasks from 1 to 3 tasks executed in parallel, number of communications from 5 to 20 communications, execution




Fig. 2. Crosstalk power penalty repartition for each signal of a N_w wavelengths receiver.

times of tasks from 5 to 10 cycles, volume of communication from 5 to 10. These task graphs are mapped on a 16 cluster architecture with a number of wavelengths varying from 1 to 16 (1, 2, 4, 8, and 16) in the ONoC.

For the example presented in this section, we consider a FSR equal to 0.8 nm and, as mentioned before, an optimal repartition of wavelengths onto this interval.




Fig. 3. Execution times comparison for min vs max wavelength allocation.

The Fig. 3 and Table I highlight the opportunity of wavelength allocation. In particular, the column “gain” of table I shows the reduction of execution time when the number for wavelengths in the waveguide increases. For example, for the Appli 1 and for 16 wavelengths in the waveguide, the possibility to allocate 16 wavelengths for each communication leads to an execution time reduction from 73 to 48, i.e. which corresponds to 34.2% of reduction. For the example “Appli 5”, which corresponds to a larger task graphs with more communications between tasks, the possibility to support

greater wavelengths allocation leads to an important reduction in terms of the execution time of the application (from 147 to 85, i.e 42% of execution time reduction).

Fig. 4 presents the limits of the Pareto fronts for the application 2 (the same type of curves can be observed for the other task graphs), and for a number of wavelengths available in the waveguide varying from 1 to 16. This figure clearly shows that increasing the number of wavelengths in a waveguide has a great impact in terms of energy penalty, due to wavelengths crosstalks between communications.

Furthermore, this figure also shows that increasing the number of wavelengths in a waveguide as no real impact in terms of execution time if WDM is used only between communications (see the line “Allocation of 1 wavelength per communication”). It means that WDM must be considered not only to parallelise communications with each other, but also to parallelise each communication.

Finally, from the design space exploration point of view, our proposal enables to extract the bounds in terms of execution times and also enables to evaluate the energy penalty induced by crosstalk between wavelengths when using WDM technique. As discussed in the introduction section, these results give a global overview of the wavelengths allocation design space, and can help the designer to select the appropriate wavelength allocation for the ONoC.

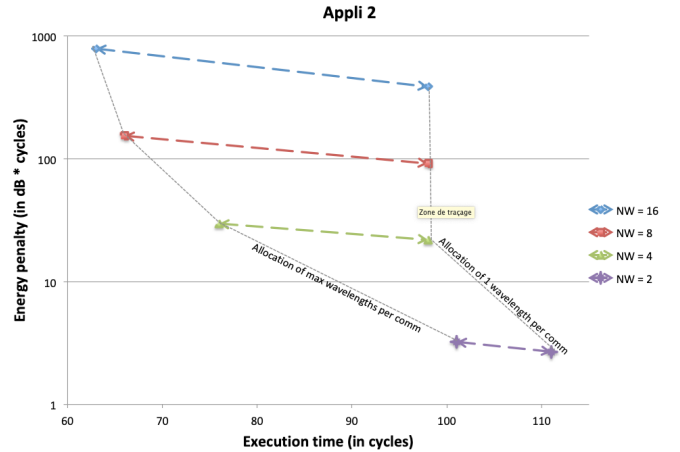


Fig. 4. Limits of Pareto fronts for application 2 (task graph 2).

V. RELATED WORKS

The wavelength allocation has been first considered as an optimisation problem where each communication message is assigned to only one wavelength [5]. The objective function and the set of constraints ensure the maximum throughput by spatially sharing the wavelengths between different communications. However, this work does not benefit from multiple wavelengths per communication to speed up applications. In [7], the authors propose a Multiple Ring-based Optical NoC and a wavelengths assignation based on WDM is used to reduce the contention during communication and to simplify the control of the NoC. In [9], the authors propose to use

TABLE I
EXECUTION TIME COMPARISONS BETWEEN MONO WAVELENGTH ALLOCATION AND MULTI WAVELENGTHS FOR EACH COMMUNICATION.

Applications	(# tasks, # comm)	1 WL				2 WL			4 WL			8 WL			16 WL		
		1 wl	1 wl	1..2 wl	gain %	1 wl	1..4 wl	gain %	1 wl	1..6 wl	gain %	1 wl	1..8 wl	gain %	1 wl	1..16 wl	gain %
Appli 1	(6, 5)	82	73	63	13.7	73	55	24.7	73	50	31.5	73	48	34.2			
Appli 2	(8, 11)	160	111	101	9.0	98	76	22.4	98	66	32.6	98	63	35.7			
Appli 3	(10, 18)	196	128	121	5.5	112	86	23.2	112	80	28.5	112	73	34.8			
Appli 4	(12, 18)	219	161	143	11.2	153	114	25.5	153	99	35.2	153	90	41.2			
Appli 5	(12, 20)	225	158	139	12.0	147	105	28.6	147	89	39.4	147	85	42.2			
Average gain		-			10.2			24.9			30.8			34.2			
Waveguide configuration:		# WL: Nb of wavelengths available on the waveguide #..# wl: range of possible wavelengths allocated for each communication															

WDM technique through deterministic routing algorithm to eliminate the need for resource reservation. The ONoC is then shared for different communications, but a pair of sender and receiver uses only one wavelength for the data transmission. In the paper [12], the authors compare the complexity of different solutions for the schedule of packets on WDM ring topology. The solution considers time-slotted allocation for the wavelengths, and different policies are considered. This work proposes to split the message to reduce the design cost.

For specific communication needs, i.e. broadband communications, the authors of [4] take advantage of WDM to support data channels on multiple parallel wavelengths, in order to distribute the information to several receivers.

In [14], the impact of crosstalk in ONoC is studied and a generic approach is defined to model crossing angles equals to 60° and 120° rather than conventional 90° crossing.

In [11], the authors have developed a genetic algorithm to explore the design space of communication management in the context of WDM technique. The technique is an evolution of the ORNoC proposed in [10]. The algorithm provides a set of large number of solutions and a Pareto Front can be extracted from this set, but due to the genetic algorithm progress, bounds of the design space can be missed. Furthermore, the time to extract the set of solutions can be very high even if the design space is small.

VI. CONCLUSION

Due to its intrinsic features, Wavelength Division Multiplexing in the ONoC domain is an interesting opportunity to offer very high bandwidth for the data transfers between cores and/or clusters of cores. While this technique is generally used to ensure several communications in parallel on the same waveguide, WDM also enables to use several wavelengths for the same communication. However, this technique induces crosstalk between wavelengths when they are used at the same time on the same part of waveguide. Then exploring the design space of wavelength allocation is a difficult task for the designer. To tackle this problem, this article presents a mathematical formulation to extract the limits of Pareto front in terms of execution time versus crosstalk energy penalty.

Two mathematical formulations are defined, one for each limit of the Pareto front, and a set of constraints ensuring the temporal scheduling of tasks and the spatial and temporal distribution of communications onto the ONoC. From this mathematical formulation, we demonstrate that our proposal

can help the designer during the design space exploration by highlighting the optimisation opportunity in terms of execution time of his application, but it also gives information about the energy overhead needed to improve the performance.

REFERENCES

- [1] T. Alexoudi, N. Terzenidis, S. Pitris, M. Moralis-Pegios, P. Maniotis, C. Vagionas, C. Mitsolidou, G. Mourgias-Alexandris, G. T. Kanellos, A. Miliou, K. Vysokinos, and N. Pleros. Optics in computing: From photonic network-on-chip to chip-to-chip interconnects and disintegrated architectures. *Journal of Lightwave Technology*, 37(2):363–379, Jan 2019.
- [2] M. Bahadori, S. Rumley, H. Jayatilaka, K. Murray, N. A. F. Jaeger, L. Chrostowski, S. Shekhar, and K. Bergman. Crosstalk penalty in microring-based silicon photonic interconnect systems. *Journal of Lightwave Technology*, 34(17):4043–4052, 2016.
- [3] L. Benini and G. De Micheli. Networks on chips: a new soc paradigm. *Computer*, 35(1):70–78, Jan 2002.
- [4] A. Biberman, B. G. Lee, N. Sherwood-Droz, M. Lipson, and K. Bergman. Broadband operation of nanophotonic router for silicon photonic networks-on-chip. *IEEE Photonics Technology Letters*, 22(12):926–928, June 2010.
- [5] I. Cerutti, N. Andriolli, P. Pintus, S. Faralli, F. Gambini, O. Liboiron-Ladouceur, and P. Castoldi. Fast scheduling based on iterative parallel wavelength matching for a multi-wavelength ring network-on-chip. In *International Conference on Optical Network Design and Modeling (ONDM)*, pages 180–185, May 2015.
- [6] Z. Chen, H. Gu, Y. Chen, and H. Zhang. Wavelength assignment in optical network-on-chip: Design and performance. In *IEEE International Conference of IEEE Region 10 (TENCON 2013)*, pages 1–4, 2013.
- [7] H. Gu, K. Chen, Y. Yang, Z. Chen, and B. Zhang. Mronoc: A low latency and energy efficient on chip optical interconnect architecture. *IEEE Photonics Journal*, 9(1):1–12, Feb 2017.
- [8] R. Ho, K. W. Mai, and M. A. Horowitz. The future of wires. *Proceedings of the IEEE*, 89(4):490–504, April 2001.
- [9] S. Koochi and S. Hessabi. All-optical wavelength-routed architecture for a power-efficient network on chip. *IEEE Transactions on Computers*, 63(3):777–792, March 2014.
- [10] S. Le Beux, J. Trajkovic, I. O’Connor, G. Nicolescu, G. Bois, and P. Paulin. Optical ring network-on-chip (ornoc): Architecture and design methodology. In *Design, Automation Test in Europe*, pages 1–6, 2011.
- [11] J. Luo, A. Elantably, V. D. Pham, C. Killian, D. Chillet, S. Le Beux, O. Sentieys, and I. O’Connor. Performance and energy aware wavelength allocation on ring-based wdm 3d optical noc. In *Design, Automation Test in Europe Conference Exhibition (DATE)*, pages 1372–1377, 2017.
- [12] B. Uscumlic, A. Gravey, I. Cerutti, P. Gravey, and M. Morvan. The impact of network design on packet scheduling in slotted wdm packet rings. In *International Conference on Photonics in Switching*, pages 1–2, Sep. 2009.
- [13] Xiaowen Wu, Jiang Xu, Yaoyao Ye, Zhehui Wang, Mahdi Nikdast, and Xuan Wang. Suor: Sectioned unidirectional optical ring for chip multiprocessor. *J. Emerg. Technol. Comput. Syst.*, 10(4), June 2014.
- [14] Y. Xie, J. Xu, J. Zhang, Z. Wu, and G. Xia. Crosstalk noise analysis and optimization in 5x5 hitless silicon-based optical router for optical networks-on-chip (onoc). *Journal of Lightwave Technology*, 30(1):198–203, 2012.