Microprocessors and Microsystems 43 (2016) 4–13

Contents lists available at ScienceDirect

Microprocessors and Microsystems

journal homepage: www.elsevier.com/locate/micpro

Minimizing power loss in optical networks-on-chip through application-specific mapping

Edoardo Fusella, Alessandro Cilardo*

Department of Electrical Engineering and Information Technologies, University of Naples Federico II, Napoli, Italy

ARTICLE INFO

Article history:
Received 1 June 2015
Revised 22 November 2015
Accepted 16 January 2016
Available online 4 April 2016

Keywords:
On-chip interconnect
System-on-chip
Embedded systems
Design automation
Silicon photonics
Optical network-on-chip
NoC

ABSTRACT

The next generation of Multiprocessor Systems-on-Chip will require communication facilities that cannot be provided by traditional electronic communication infrastructures. Silicon photonics appears today a promising solution to handle future communication needs thanks to ultra-high bandwidth and extremely low-power consumption. However, designing an optical on-chip network requires addressing several challenges that have no equivalent in the electronic domain. In particular, photonic networks-on-chip suffer from considerable power loss, which affects the network scalability and impacts the performance by constraining the total number of wavelengths and hence the available bandwidth. In this paper, we propose an algorithm which automatically maps the IP cores onto a generic mesh-based photonic NoC architecture such that the worst-case power loss is minimized. As the main contribution, we first formulate the problem of power loss aware mapping and then we propose a genetic algorithm to solve it. Experimental results show that the power loss can be significantly reduced enabling much higher scalability of the on-chip photonic interconnect.

1. Introduction

The restless grow in transistor integration has pushed the semiconductor industry to a shift from the single-core era to a multicore paradigm during the last decades. In a large-scale multicore scenario, an energy-efficient on-chip communication fabric is a key aspect ensuring performance scalability. While traditional electronic interconnects are constrained by physical limitations in terms of power dissipation, latency, and bandwidth [1], silicon photonics appears a promising path to energy-efficient ultra-high bandwidth on-chip communication. Nanophotonic waveguides, the photonic counterpart of a wire, can in fact achieve bandwidths in the order of terabits per second by exploiting wavelength division multiplexing (WDM), while photonic signaling consumes less power than electrical interconnects [2].

In order to exploit the potential advantages of photonics, it is necessary to deal with an intrinsic characteristic of photonic devices, the insertion loss, i.e. the power loss that a photonic element induces when it is inserted in an optical path. The insertion loss is one of the major limitations when designing a photonic NoC. In fact, the power of an optical signal must be above a certain threshold when arriving at the photodetectors in order to ensure a proper detection. As a consequence, the power injected into the chip must be higher than the photodetector sensitivity plus the worst case power loss. However, the total power cannot exceed a certain threshold due to the nonlinearities of the silicon material. In case of multiwavelength signals, this problem is exacerbated since these considerations must apply to each individual wavelength channel. Under the assumption of equal power for each wavelength channel, the above constraints yield the following inequality [3]:

$$P_{Budget}^{dB} \geq IL_{wc}^{dB} + 10\log_{10} N_{\lambda} \quad (1)$$

where P_{Budget}^{dB} is the optical power budget, representing the difference between the upper limit in transmission power and the photodetector sensitivity, IL_{wc} is the worst-case insertion loss, and N_{λ} is the total number of wavelength channels.

Power loss is a major goal when designing a photonic NoC architecture, since a high power loss may easily result in a network with poor performance or even an inoperable architecture. As a major insight of this work, we recognize that mapping optimization based on specific communication patterns is the best chance to face this problem in application-specific multi-core systems-on-chip. In fact, the NoC architecture can be customized for a target application when its traffic characteristics are known at design time, which is the case for most embedded applications running on multiprocessor Systems-on-Chip (MPSoCs). The problem of mapping a set of given IP cores to the NoC tiles is

* Corresponding author. Tel.: +39 0817683852; fax: +39 0817683816.
E-mail address: acilardo@unina.it (A. Cilardo).

http://dx.doi.org/10.1016/j.micpro.2016.01.014
0141-9331/

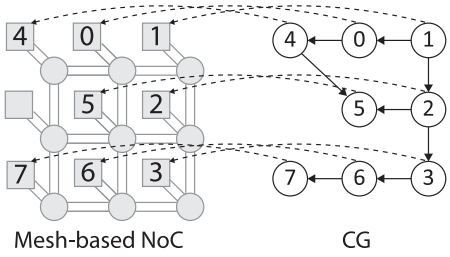

Fig. 1. A mesh-based on-chip architecture and an example of mapping problem.

illustrated in Fig. 1. An application, previously divided into a graph of concurrent tasks, should be assigned to a set of available cores by mapping different functions to different regions of the system, so that the traffic exhibits power loss optimized patterns.

In the above scenario, this paper addresses for the first time the problem of power-loss aware mapping for photonic NoCs. The work introduces a framework for the automated design space exploration of mapping solutions for photonic networks, based on a suitable formalization of the optimization problem, described in the paper, as well as a genetic algorithm to solve it. The experimental results show that the power loss can be significantly reduced, enabling much higher scalability of the on-chip photonic interconnect.

The paper is organized as follows. Section 2 summarizes the main related contributions in the technical literature. Section 3 describes the tile-based optical NoC architecture and the related power loss model adopted in this work. Section 4 defines the optimization problem addressed by this work; the section also describes the genetic algorithm developed to solve the problem. Section 5 discusses the results achieved by presenting the experimental setup and several real-world applications. Section 6 gives some final remarks.

2. Background

Optical interconnects are widely used to provide communication facilities over long and medium distances. However, thanks to recent progress in nanophotonic technologies, silicon photonics also appears today a promising approach for on-chip communication. Photonic networks proposed in the literature can be mainly divided into two classes, passive or active according to the type of the Microring Resonator (MR) used. Usually passive networks are wavelength-routed optical networks (WRONoCs), while active networks are circuit-switched networks indicating the use of wavelength selectivity in order to implement respectively routing facilities or bandwidth aggregation.

In case of WRONoCs, routing facilities can be obtained by statically or dynamically assigning a wavelength to a node or to a pair of nodes. A communication between two nodes can start by simply using the right wavelength channel through MR modulators and filters. Wavelength routing can enable extremely short latencies since, with no conflicts, the propagation delay is simply the time of flight at the speed of light. The downside is that the available optical spectrum is extremely reduced and scaling the network size is limited by the maximum number of wavelengths available.

A large body of work addresses WRONoCs. In [4,5] a clustered architecture with local electronic buses and a global optical ring connecting local domains is proposed. The optical configuration requires that each wavelength must be statically assigned to a given optical network interface, limiting the achievable parallelism. In [6], the Corona Multiple-Writer Single-Reader (MWSR) photonic ring is presented. The main drawback of this architecture consists of the high implementation complexity in terms of number of microring resonators required. In addition, Corona requires arbitration to handle write conflicts leading to arbitration overheads. Pan et al., [7] proposes Firefly, a clustered architecture locally connected via an electronic medium and globally via an optical crossbar partitioned into multiple logical crossbars. The waveguides are configured in a reservation-assisted SWMR configuration, where dedicated reservation channels are used to highly reduce the power consumption. On the other hand, the reservation process and token round-trip overhead cause additional latency. Pan et al. [8] extends [7] by enabling the generation of a new token each cycle and hence reducing the latency overhead of the token arbitration. Kirman and Martnez [9] proposes an all-optical torus network with a wavelength-based oblivious routing where source-destination pairs do not require a static assigned wavelength, relying on a wavelength reuse policy. While this allows a drastic reduction of the number of required wavelengths, the architecture is still subjected to scalability issues. In [10], the ORNoC architecture is presented. ORNoC is a contention-free ring-based on-chip network that uses an automatic wavelength/waveguide assignment. Bahirat and Pasricha [11] presents METEOR, a complete communication architecture consisting of a reconfigurable electrical mesh coupled with an optical ring whose waveguides could be configured as a combination of SWMRs and MWMRs. Le Beux et al. [12] introduces CHAMELEON, a SWSR optical network-on-chip able to exploit reconfigurability both at run- or compile-time in order to open and close dedicated channels between cores.

Unlike wavelength-routed optical networks, when the wavelength selectivity is used with a bandwidth aggregation purpose, high-bandwidth channels are obtained by multiplexing data onto many parallel wavelengths. These architectures usually combine electronic and photonic technologies to build a hybrid network made up of two subnetworks: an electronic packet-switched network for handling control messages and a photonic circuit-switched network for data messages [2]. A hybrid torus-based NoC made up of a photonic circuit-switched and an electronic packet-switched network was presented in [2]. The main drawback concerns the power loss which affects the architecture due to the large number of waveguide crossings in the topology layout. In [13], the authors propose HOME, a hybrid optical-electronic mesh-based NoC exploiting hybrid optical-electronic routers able to perform a selective transmission. In order to reduce the energy consumption and achieve a higher bandwidth, Chan et al. [14] first performs a topology comparison between different architectures and then presents an improved topology for the architecture presented in [2]. Differently, [15] proposes a hybrid wireless/optical-wired architecture where the control messages are broadcast via the photonic antennas. In [16] the THOE architecture is proposed. THOE is a low cost torus-based hybrid optical-electronic NoC that can be implemented with a lower number of waveguides and ring resonators compared to [2].

Despite its importance, the research on power-loss aware architecture design has never been covered extensively. In that respect, a few works address the power loss aware design of the photonic switching element (PSE) where microring resonators[1] are added at waveguide crossings to provide switching facilities. Shacham et al. [2] first described a 4 × 4 photonic basic switch, while [17] proposed an optimized version providing less insertion loss. Gu et al. [18] presented a 5 × 5 switch whose internal structure of the crossbar is restructured to reduce the number of waveguide crossings and hence the whole insertion loss. In [19], *OXY* is presented. OXY is a 5 × 5 switch whose physical layout is optimized to reduce its insertion loss and crosstalk noise. While the above

[1] A microring resonator is a waveguide forming a closed loop which can be designed to perform switching functions: it injects a '0' value when coupling the light corresponding to its resonant wavelength, while it injects a '1' value when letting pass the light.