CONCURRENCY AND COMPUTATION: PRACTICE AND EXPERIENCE
Concurrency Computat.: Pract. Exper. (2014)
Published online in Wiley Online Library (wileyonlinelibrary.com). DOI: 10.1002/cpe.3331

EDITORIAL

Exploiting silicon photonics for energy-efficient heterogeneous parallel architectures

Welcome to this special issue of the journal Concurrency and Computation: Practice and Experience on Exploiting Silicon Photonics for Energy-Efficient Heterogeneous Parallel Architectures, which contains five original manuscripts that cover a complete range of perspectives.

Silicon photonics is undoubtedly expected to play a big role in the evolution in board, cross-chip, interposer-level and on-chip interconnection for low-power and/or high-performance computer systems spanning from high-end embedded devices (e.g., tablets and smartphones) and other System-on-Chips (SoCs), up to chips for the High Performance Computing (HPC) domain.

The unique features of photonics (e.g., extreme low-latency, end-to-end transmission, high bandwidth density and passive long-range propagation) have the potential constitute a discontinuity element able to modify the expected shape of future computer systems from the design point of view and also from the programmability and/or runtime management perspectives. Summarizing, silicon photonics can bring innovations and benefits into current and foreseeable computing systems directly, due to their intrinsic features, but also indirectly enabling the evolution toward architectures, runtime and resource management approaches that maximize the photonic raw technological opportunities and lead to more efficient overall designs, otherwise impossible.

For instance, the extreme low transmission latency (i.e., group velocity of light into silicon, about 15 ps/mm) can potentially allow a higher number of architectural modules to be *close* each other and thus to enable their effective tight cooperation and communication. However, computer architecture, as well as network on- and off-chip, designs needs to be adapted to extract maximum benefits from the photonic technology, which exposes other substantial differences compared to what designers are well accustomed to. For example, at the moment optical interconnection is *end-to-end* by nature therefore much of the knowledge and solutions based on store-and-forward paradigm cannot be directly transferred and exploited. However, propagation into a silicon waveguide can occur with limited losses (e.g., even less than 1 dB/cm) over on-chip or interposer distances without signal regeneration needs. In brief, in this arena, new ad-hoc solutions need to be pursued.

Then, despite optical communication is very well established and all the involved elements (such as modulators, detectors, waveguides and resonators) have been extensively researched on, silicon photonics applied to computing systems is still in its infancy. Consequently, researchers have already highlighted a deep interaction between design choices at very different layers of abstraction. For this reason, nowadays the whole spectrum of layers, from physical concerns about optical structures on silicon (e.g., module layout and modeling to expose interactions and, for instance, to evaluate and limit insertion losses) up to network issues (e.g., connectivity and topologies, bandwidth and latency) and even computer architecture choices (e.g., memory coherency and consistency models, memory hierarchy and parallelism management), need to be studied with a strong multidisciplinary approach.

This special issue contributes to this promising field with extended and carefully reviewed versions of selected papers from the First International Workshop on Exploiting Silicon Photonics for Energy-Efficient Heterogeneous Parallel Architectures (SiPhotonics'14), which was held in Vienna (Austria) as part of the 9th HiPEAC conference on High Performance and Embedded Architecture and Compilers. Therefore, due to the peculiarity of the depicted scenario, the papers of this number address a complete range of perspectives to silicon photonics applied to computing systems, from raw technology issues and solutions up to studies at the overall system level of modern multi-/many-core systems, both from academic and industrial researchers working in this area.

We start this special issue with the paper entitled *Optical Crossbars on Chip, A Comparative Study based on Worst-Case Losses*. In this paper, Le Beux *et al.* [1] study the worst-case losses for possible crossbar implementations depending on three key design factors: network topology, considered layout and insertion losses induced by the fabrication process. They compare different implementations relying on matrix, multistage and ring-based network topologies, finding that ring-based networks yield the most power-efficient solution.

The paper *Capturing the Sensitivity of Optical Network Quality Metrics to its Network Interface Parameters* by Ortin *et al.* [2] addresses the network interface architecture (NI) required to support optical communications on the silicon chip. The paper proposes a complete network interface architecture for wavelength-routed optical NoCs, by coping with the intricacy of some specific issues such as flow control, buffering strategy and dual-clock domains, deadlock avoidance, serialization, and above all, the co-design around the requirements of a cache coherency protocol. The most important conclusion is that NI design and optimization perhaps has now higher priority over the relentless search for improvements in individual optical devices.

The emerging of circuit-level simulators for photonic integrated circuits (PICs) is driven by recent developments in technologies for integration of large-scale monolithic PICs in both, silicon and InP technologies. Arellano *et al.* [3] present their solution for modeling PICs in the framework of the circuit-level simulation tool *VPIcomponentMaker^{TM} Photonic Circuits*. In their paper *The Power of Circuit Simulations for Designing Photonic Integrated Circuits*, they demonstrate the combination of different simulation approaches in time domain, frequency domain and time-and-frequency domain (TFDM) for fast and accurate simulations. This is particularly crucial for being able to model and design more and more complex optical circuits like the ones that could be needed to be employed in, and/or between, chip multiprocessors.

In the paper entitled *Managing Resources Dynamically in Hybrid Photonic-Electronic NoCs*, García-Guirado *et al.* [4] present novel fine-grain policies to manage the photonic resources in a tiled-CMP scenario. The objective is to maintain the optical channel in the load condition that allows it to deliver best performance. Their policies are dynamic and base their decisions on parameters such as message size, ring availability and distance between endpoints, at the message level. The resulting network behavior is also fairer to all cores, reducing processor idle time thanks to faster thread synchronization, improving performance and reducing both the overall network latency and energy consumption when compared to the same CMP without the photonic ring.

Finally, the paper *Towards Zero Latency Photonic Switching in Shared Memory Networks* explores techniques which intelligently use information from the memory hierarchy to predict communication in order to setup photonic circuits with reduced or eliminated arbitration latency in case of reconcilable optical networks. Madarbux *et al.* [5] present a switch scheduling algorithm which arbitrates on a per memory transaction basis and holds open photonic circuits to exploit temporal locality, showing that this can reduce the average arbitration latency overhead and eliminate arbitration latency altogether for many of memory transactions.

ACKNOWLEDGEMENTS

We would like to thank all the authors, reviewers and editors involved in the elaboration of this special issue, including also the reviewers that were involved in the SiPhotonics'14 workshop, where short versions of the papers were previously selected. We are especially grateful to Profs. Geoffrey C. Fox and David W. Walker, editors of the journal, for approving this special issue and for his help along the process of its preparation.

REFERENCES

1. Le Beux S, Li H, Nicolescu G, Trajkovic J, O'Connor I. Optical crossbars on chip, A comparative study based on worst-case losses. *Concurrency and Computation: Practice and Experience* 2014; DOI: 10.1002/cpe.3336.
2. Ortín-Obón M, Ramini L, Viñals V, Bertozzi D. Capturing the sensitivity of optical network quality metrics to its network interface parameters. *Concurrency and Computation: Practice and Experience* 2014; DOI: 10.1002/cpe.3330.
3. Arellano C, Mingaleev SF, Sokolov ES, Richter A. The power of circuit simulations for designing photonic integrated circuits. *Concurrency and Computation: Practice and Experience* 2014; DOI: 10.1002/cpe.3335.

4. García-Guirado A, Fernández-Pascual R, García JM, Bartolini S. Managing resources dynamically in hybrid photonic-electronic NoCs. *Concurrency and Computation: Practice and Experience* 2014; DOI: 10.1002/cpe.3332.
5. Ridwan Madarbux M, Van Laer A, Watts PM, Jones TM. Towards zero latency photonic switching in shared memory networks. *Concurrency and Computation: Practice and Experience* 2014; DOI: 10.1002/cpe.3334.

SANDRO BARTOLINI
Dipartimento di Ingegneria dell'Informazione e Scienze Matematiche,
University of Siena, Italy
E-mail: bartolini@dii.unisi.it

JOSÉ M. GARCÍA
Departamento de Ingeniería y Tecnología de Computadores,
University of Murcia, Spain
E-mail: jmgarcia@ditec.um.es

DOI: 10.1002/cpe