Guest Editors' Introduction: Emerging Networks-on-Chip Designs, Technologies, and Applications

EDOARDO FUSELLA, University of Naples Federico II

MAHDI NIKDAST, Colorado State University

IAN O'CONNOR, Ecole Centrale de Lyon

JOSÉ FLICH, Universitat Politècnica de València

SUDEEP PASRICHA, Colorado State University

ACM Reference format:

Edoardo Fusella, Mahdi Nikdast, O'Connor Ian, José Flich, and Sudeep Pasricha. 2019. Guest Editors' Introduction: Emerging Networks-on-Chip Designs, Technologies, and Applications. *J. Emerg. Technol. Comput. Syst.* 15, 1, Article 1 (February 2019), 2 pages.

<https://doi.org/10.1145/3296021>

Moore's Law has enabled a continuous growth in transistor integration for decades, pushing the semiconductor industry to move from the era of single-processor frequency scaling to the era of parallelization. Since the introduction of multicore architectures and Chip Multiprocessors (CMPs), it quickly became clear that on-chip communication would play a major role in the overall performance and cost of the system. Performance, energy consumption, power dissipation, thermal aspects, predictability, reliability, and fault tolerance represent just a brief list of the metrics affected by on-chip communication. On these grounds, the Network-on-Chip (NoC) paradigm has emerged as a solution to face the complex communication problems that arise as the number of on-chip components increases. Compared with the traditional global interconnects (e.g., bus-based solutions), NoCs are based on a modular packet-switched architecture, providing a scalable interconnection fabric with high-bandwidth, low-latency, energy-efficient communication while satisfying tight area and power budgets. Nowadays, the current ultra-high level of integration (100s of millions of transistors) has led to the widespread adoption of the NoC paradigm. Furthermore, as the on-chip integration density continues to increase to satisfy emerging applications demands (e.g., big data analytics and machine learning), novel interconnect technologies based on silicon photonics and wireless are emerging to improve the communication performance in such large-scale CMPs.

Authors' addresses: E. Fusella, Department of Electrical Engineering and Information Technologies, University of Naples Federico II, via Claudio 21, 80125 Naples, Italy; email: edoardo.fusella@unina.it; M. Nikdast and S. Pasricha, Department of Electrical and Computer Engineering, Colorado State University, Fort Collins, CO 80523-1373, USA; emails: {mahdi.nikdast, sudeep.pasricha}@colostate.edu; I. O'Connor, Ecole Centrale de Lyon, INL, Ecully, F-69134, France; email: ian.oconnor@ec-lyon.fr; J. Flich, Department of Computer Engineering (DISCA), Universitat Politècnica de València, Valencia, Spain, 46071; email: jflich@disca.upv.es.

Permission to make digital or hard copies of part or all of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for third-party components of this work must be honored. For all other uses, contact the owner/author(s).

© 2019 Copyright held by the owner/author(s).

1550-4832/2019/02-ART1

<https://doi.org/10.1145/3296021>

Nevertheless, several challenges still remain to be addressed to exploit the full potential of these interconnect technologies at any abstraction level, from the very bottom level of device manufacturing to the architecture and system level. While NoCs have evolved from fairly simple communication fabrics to massively complex interconnect systems, a major effort is still required for harnessing the full potential of emerging NoCs in future computing scenarios. The purpose of this special issue was to solicit papers that described the latest advances in the design and technology for emerging NoC architectures and evaluate their future prospects.

This special issue received 16 original submissions, and we assigned each manuscript to at least three reviewers. After the multiround peer-review process, the following seven articles were accepted for publication: (1) “Design and Multi-Abstraction Level Evaluation of a NoC Router for Mixed-Criticality Real-Time Systems,” (2) “Time-Randomized Wormhole NoCs for Critical Applications,” (3) “Limit of Hardware Solutions for Self-Protecting Fault-Tolerant NoCs,” (4) “Fault-Tolerant Network-on-Chip Design with Flexible Spare Core Placement,” (5) “Thermal-Aware Test Scheduling Strategy for Network-on-Chip-Based Systems,” (6) “PANE: Pluggable Asynchronous Network-on-Chip Simulator,” and (7) “BigBus: A Scalable Optical Interconnect.”

The first two articles address the demands of critical real-time applications in NoCs. In “[Design and Multi-Abstraction Level Evaluation of a NoC Router for Mixed-Criticality Real-Time Systems](#),” Dridi et al. present the Double Arbiter and Switching (DAS) router, a novel NoC router design to support mixed-criticality systems based on an accurate worst-case communication time analysis for high-critical flows. In the next article, “[Time-Randomized Wormhole NoCs for Critical Applications](#),” Slijepcevic et al. propose a wormhole-based NoC design able to provide better performance guarantees than deterministic approaches by making use of a wormhole router with randomized arbitration.

Articles three and four present two techniques for increased reliability and fault tolerance in NoCs. In particular, Louri et al., in “[Limit of Hardware Solutions for Self-Protecting Fault-Tolerant NoCs](#),” study the limits of hardware solutions for self-protecting NoCs by considering Self-Test and Self-Protect (STAP) architectures as well as Triple Modular Redundancy (TMR) with voting. In addition, Bhanu et al., in the fourth article, titled “[Fault-Tolerant Network-on-Chip Design with Flexible Spare Core Placement](#),” propose a spare core placement technique for mesh-based NoCs providing an improved fault tolerance in the network.

In the next article, “[Thermal-Aware Test Scheduling Strategy for Network-on-Chip-Based Systems](#),” Kanchan et al. present a preemptive test scheduling technique to improve the test time by reducing resource conflict and a thermal-aware test scheduling technique to test cores in 2D as well as 3D stacked NoC-based multicore systems. Following this, Ved et al. in “[PANE: Pluggable Asynchronous Network-on-Chip Simulator](#)” propose PANE, a NoC simulator that allows system-level simulation of asynchronous NoCs.

In the last article, “[BigBus: A Scalable Optical Interconnect](#),” Bashir et al. introduce BigBus, a novel optical communication architecture aimed at enhancing performance while reducing power consumption by means of a combination of laser modulation and power sharing across optical transceivers in emerging NoCs.

We believe the NoC topic will continue to be of growing interest and importance for both academia and industry. We thank all those who contributed to this special issue. In particular, we are grateful to the authors for their contributions and to the reviewers who helped the authors enhance the quality of their manuscripts. We would also like to thank the ACM JETC Editor-in-Chief Prof. Yuan Xie and the Journal Administrator G. Codina for invaluable guidance and assistance while putting together this special issue.