Guest Editorial
New Interconnect Technologies in On-Chip Communication

WITH Moore's Law continuing to increase the number of transistors, the number of cores on a single chip is continuing to increase as well. The International Technology Roadmap for Semiconductors (ITRS) projects that device scaling will continue but the interconnect will not scale at the same rate. As a result, communication or movement of data is becoming the bottleneck of manycore systems and often determines the scalability, both in terms of cost and performance. To overcome the limitations of conventional electrical signaling communication, alternative technologies have been recently proposed and will not only impact how data is communicated on-chip but also the system design. The on-chip communication or the network-on-chip architecture (NoC) needs to be designed efficiently to properly exploit these new interconnect technologies to enable the additional computation capability of manycore processors to be properly leveraged.

The purpose of this special issue is to report on recent advances in new interconnect technologies and the design of on-chip communication architecture and circuits that leverage these new technologies. In particular, this special issue focuses on enabling circuit and system designs that exploit these new technologies to create cost-efficient communication in manycore processors. Recent advances in devices have shown the feasibility of some alternative technologies for on-chip communication. However, it still remains an open problem in terms of how these individual components can be scaled through different circuit and architecture designs to create a system that fully leverages the benefit of these technologies while minimizing the overall cost of the system. New interconnect technologies often introduce additional cost in terms of area and/or power consumption and the trade-off from these technologies needs to be properly explored, compared with the state-of-the-art advances in conventional, electrical signaling. Some technologies while providing significant benefits (e.g., lower latency and higher bandwidth) can introduce additional challenges that did not exist with conventional electrical signaling, such as reliability or thermal issues.

The advances in different interconnect technologies present different challenges, which include determining the appropriate architectures (both microarchitecture and system architecture) that can properly exploit these new technologies. The different technologies can lead to different optimal topology for the network-on-chip, compared with conventional electrical signaling. The new interconnect technologies can also enable communication organization that was not feasible with electrical signaling—such as providing a scalable, global bus or providing a fully-connected network. Some of the technologies can be combined to further increase the benefits—for example, 3-D stacking is often combined with other interconnect technologies as the additional layer is leveraged for the other interconnect technology. The challenges are not only in the on-chip communication but also in off-chip communication with main memory. As the number of cores increases, not only is higher memory bandwidth required to supply data to all the cores but the chip-to-chip communication power can consume a major portion of the total system power.

The papers in this special issue address these challenges from both the circuit and architecture perspective for various interconnect technologies. This special issue received 83 submissions and among them, the Guest Editors preselected 51 papers which went through a review process. Many of the papers went through a revision process and in the end, 17 papers were accepted by the reviewers and the Guest Editors for publication in this special issue. The focus of this journal was not on traditional network-on-chip papers but focused on *technology-aware* network-on-chip papers. As a result, there were several strong submissions that we decided not to include in the final issue as their focus was more on traditional network-on-chip microarchitecture. Although we attempted to cover different interconnect technologies, this special issue is not meant to be a comprehensive coverage.

The first paper is an introductory paper from the Guest Editors that provides an overview of alternative interconnect technologies and their impact on on-chip network architecture. The following three papers in this issue describe how nanophotonic technology can be leveraged to provide efficient global communication. Batten *et al.* provide a summary of nanophotonics technologies and survey different architectures leveraging nanophotonics. They also present different cases studies of using nanophotonics which include on-chip network, processor-to-DRAM network, and memory channel design. Wassel *et al.* present nanophotonics NoC architectures that leverage plasmonic components to avoid manufacturing-induced variability and the high degree of temperature dependence in micro-ring resonators. Nitta *et al.* combine nanophotonics with 3-D technology to create a fully connected topology and avoid having to globally arbitrate for the optical channels.

The next set of papers discuss how transmission lines can be used to enable alternative communication architecture. Carpenter *et al.* present how transmission lines can be used for global on-chip communication and avoid packet switching by leveraging transmission lines to create a globally shared bus for medium-scale chip multiprocessors. Hsu *et al.* present an alternative implementation scheme for on-chip transmission

Digital Object Identifier 10.1109/JETCAS.2012.2196890

lines that provide better attenuation characteristics and higher bandwidth compared to conventional on-chip transmission line designs by twice-interleaved metal patterns in lower-level metal layers. Using transmission lines combined with off-chip multi-band radio-frequency (RF) signaling, Therdsteerasukdi *et al.* present a memory interconnect architecture that leverages a hierarchical tree-based topology to improve scalability and overall system throughput in a many-DIMM DRAM system. Kim *et al.* propose a noncoherent amplitude shift keying modulation-based RF-interconnect memory interface to increase interconnect bandwidth. In addition to RF, wireless is an another alternative technology for global on-chip communication and Deb *et al.* discuss the challenges and opportunities with wireless interconnect technology.

One of the critical technology to enable 3-D ICs is how the different dies communicate with each other. Commonly used method is the through-silicon vias and Kim *et al.* investigate the impact of TSVs on the quality of 3-D ICs for current and future process technologies. Instead of using TSVs, Miura *et al.* present the design of an inductive-coupling interface between GPU and DRAM chip stacked together to provide high bandwidth at low energy. In comparison, Myat Thu Linn Aung *et al.* propose a novel interconnect using a capacitatively-coupled interface that allows simultaneous, bi-directional signaling using a single set of die-to-die capacitors.

The following papers exploits 3-D IC technology in different ways. Nandakumar and Marek-Sadowska takes advantage of vertical slit field effect transistors in a 3-D IC to design a more efficient hybrid NoC architecture. By only implementing the NoC on the VeSFET layer, the inter-core network of a multi-core processor can be densely built without requiring additional area on the core layers, resulting in an overall denser system. Creating a high-radix crossbar switch has often been assumed infeasible compared with multi-hop NoC but Sewell *et al.* challenge this assumption by proposing a swizzle-switch crossbar to create a scalable crossbar for both 2-D and 3-D manycore systems. Beanato *et al.* propose a modular design where one or more identical chips, each equipped with multiple cores and a network-on-chip can be stacked to provide the manufacturer with a range of system capabilities with redesigning the chips or the overall system.

The next two papers discusses circuit designs to improve data rate efficiently. Ho *et al.* propose a robust CMOS repeater design with a high-boosting pre-driver for subthreshold interconnect design. They demonstrate better trade-off between data rate and energy under supply voltages of 0.15 V and lower, compared to conventional approaches. Lee *et al.* present an energy-efficient high frequency compensated crosstalk and ISI equalizer for multi-channel, capacitively driven on-chip interconnects. They compensate crosstalk at high frequencies by inserting a zero in the aggressor signal path and achieves higher data rate. The last paper by Zhang *et al.* presents an approach to reducing power consumption of interconnect through exponential wire shaping utilizing the optical proximity correction process while keeping the wire delay unchanged.

We would like to thank all the authors who submitted to this special issue as well as the reviewers who made this special issue possible. We would also like to thank the Editor-in-Chief of JETCAS and JETCAS board for giving us this opportunity to organize this special issue. We also would like to express our deepest gratitude to the CASS personnel for their support through out the entire process. Through this issue, we hope that the readers will obtain an understanding of contemporary research in the area of new interconnect technology and their impact on circuit and system, and provide opportunities for further research and innovation in this field.

KIYOUNG CHOI, *Guest Editor*
School of Electrical Engineering and Computer Science
Seoul National University
Seoul, 151-742 Korea

JOHN KIM, *Guest Editor*
Department of Computer Science
Korea Advanced Institute of Science and Technology
Daejeon, 305-701 Korea

GABRIEL LOH, *Guest Editor*
Advanced Micro Devices, Inc.
Bellevue, WA 98007 USA

Kiyoung Choi (M'88–SM'08) received the B.S. degree in electronics engineering from Seoul National University, Seoul, Korea, in 1978, the M.S. degree in electrical and electronics engineering from Korea Advanced Institute of Science and Technology, Seoul, Korea, in 1980, and the Ph.D. degree in electrical engineering from Stanford University, Stanford, CA, in 1989.

From 1989 to 1991, he was with Cadence Design Systems, Inc. In 1991, he joined the faculty of the School of Electrical Engineering and Computer Science, Seoul National University. His primary research interests include various aspects of computer-aided electronic systems design including embedded systems design, high-level synthesis, and low-power systems design. He is also interested in computer architecture and especially in configurable and reconfigurable architecture design. He was an Associate Editor for the ACM *Transactions on Design Automation of Electronic Systems* (2003–2008). He was a Guest Editor for the ACM *Transactions on Embedded Computing Systems* (2008).

Dr. Choi is currently on the editorial board of the IEEE JOURNAL ON EMERGING AND SELECTED TOPICS IN CIRCUITS AND SYSTEMS. He was Technical Program Co-Chair (2003) and General Co-Chair (2004) of ACM/IEEE

International Symposium on Low Power Electronics and Design (ISLPED), Technical Program Co-Chair (2006) and General Co-Chair (2007) of the International Conference on Hardware/Software Codesign and System Synthesis (CODES+ISSS), and Technical Program Vice Co-Chair (2007) and Technical Program Co-Chair (2008) of the Asia and South Pacific Design Automation Conference (ASP-DAC).

John Kim (M'09) received the B.S. and M.Eng. degrees in electrical engineering from Cornell University, Ithaca, NY, in 1997 and 1998, respectively, and the Ph.D. degrees in electrical engineering from Stanford University, Stanford, CA, in 2008.

He is currently an Assistant Professor in the Department of Computer Science, Korea Advanced Institute of Science and Technology (KAIST), with joint appointment in the Web Science and Technology Division at KAIST. He spent several years working on the design of different microprocessors at Motorola and Intel. His research interests includes multicore architecture, interconnection networks, and datacenter architecture.

Dr. Kim is a member of the ACM.

Gabriel Loh (M'05–SM'09) Gabe received the B.E. degree in electrical engineering from the Cooper Union, New York, and the M.S. and Ph.D. degrees and the M.S. degree in computer science from Yale University, New Haven, CT.

He is a Principal Researcher at Advanced Micro Devices (AMD). He was also a tenured Associate Professor in the College of Computing, Georgia Institute of Technology, a visiting researcher at Microsoft Research, and a Senior Researcher at Intel Corporation. His research interests include computer architecture, processor microarchitecture, emerging technologies and 3-D die stacking.

Dr. Loh is a member of the ACM.