Special Issue on Emerging System Interconnects

John Kim, *Korea Advanced Institute of Science and Technology, Daejon, 34141, South Korea*
Nam Sung Kim, *University of Illinois, Urbana, IL, 61801, USA*

Moore's law scaling has approached the fundamental limit of physics and its end is in sight. Consequently, integrating more transistors and system components in a single monolithic die has become cost ineffective and inefficient. This has led to two innovations in: 1) integration and packaging technologies that integrate multiple system components in separate dies into a single chip package; and 2) heterogeneous computing, which allows us to continue to scale performance and energy efficiency with various accelerators. In parallel, with proliferation of cloud computing enabled by large-scale datacenters, efficient use of compute, memory and storage resources through disaggregation has become very essential for cost-effective operation of the datacenters. These changes have resulted in the emergence of *system interconnect technologies*—chip-to-chip (or die-to-die) interconnection networks that efficiently connect dies within a chip package or between chips in a system board and across boards to cost-effectively scale out the system. These system interconnect technologies present unique challenges and opportunities compared to large-scale interconnection networks or network-on-chip. This Special Issue on Emerging System Interconnects addresses recent advances in system interconnect technologies, with special emphasis on Compute Express Link (CXL) as it becomes more widely adopted in the industry. In particular, the special issue comprises nine articles, with six of the articles coming from the industry. Most of the articles address the opportunities from the emerging CXL interconnect but the special issue also includes work on Universal Chiplet Interconnect express (UCIe) interconnect technologies and promising use cases for future computing systems as well as photonic interconnect for wafer-scale systems.

THIS SPECIAL ISSUE ON EMERGING SYSTEM INTERCONNECTS ADDRESSES RECENT ADVANCES IN SYSTEM INTERCONNECT TECHNOLOGIES, WITH SPECIAL EMPHASIS ON COMPUTE EXPRESS LINK (CXL) AS IT BECOMES MORE WIDELY ADOPTED IN THE INDUSTRY.

CXL TECHNOLOGY AND USE CASES

In the first article in this issue [A1], Das Sharma, who has led the standardization effort of CXL at Intel, overviews the CXL standard. This article describes Intel's CXL implementation, provides detailed performance characteristics, and proposes a composable, scaleout rack scale architecture based on CXL. In the second article [A2], Kim et al. from Samsung present one of the most anticipated use case of CXL where the authors exploit CXL to expand both the bandwidth and the capacity of the memory system for datacenter servers through software-defined memory tiering.

CXL MEMORY POOLING

The next three articles explore another much-anticipated CXL use case, memory pooling, or disaggregation for public cloud servers. In [A3], Berger et al. from Microsoft explore the design space of CXL-based memory systems that not only considers the size, reach, and topology of the memory pool for public cloud servers, but also navigates complex design constraints around performance, virtualization, and management. In [A4], Ha et al. from SK Hynix present a memory pooling system prototype that provides support for dynamic allocation or release of memory

0272-1732 © 2023 IEEE
Digital Object Identifier 10.1109/MM.2023.3240453
Date of current version 13 March 2023.

space in the pool based on memory demands of individual servers. In [A5], Gouk et al. also demonstrate a CXL-based memory pooling prototype by providing the design of a CXL switch for memory pools, an implementation of software runtime to manage the memory pools in detail, along with detailed evaluations and comparisons with a remote direct memory access (RDMA)-based implementation for key-value-store workloads.

CXL OPPORTUNITIES

CXL presents new opportunities beyond simply interconnecting components together. In [A6], Boles et al. from Micron and IBM introduce memory functions, such as enhanced support for crash-consistent transactions on persistent memory, logging and tracing, data movement, and security in their FPGA-based CXL memory controller. In [A7], Kwon et al. propose TrainingCXL that integrated both a central processing unit and persistent memory with CXL and removes the overhead of storing periodic checkpoints from the critical path of training recommendation models.

THE IMPACT OF COMMUNICATION AND INTERCONNECT WILL ONLY CONTINUE TO GROW AS SYSTEM INTEGRATION AND SCALABILITY INCREASES, AND WE HOPE THESE ARTICLES ENCOURAGE NEW RESEARCH IN THIS EXCITING AREA.

INTERCONNECT FOR CHIPLETS

The next two articles address system interconnect challenges when integrating chiplets. In [A8], Das Sharma (Intel) introduces UCIe interconnect that is an open industry standard interconnect that allows chiplets from any supplier to be packaged in an interoperable fashion. This article delves into the architectural and protocol aspects of the UCIe specification, and then presents results based on Intel's implementation. Another approach to integrating chiplets is through wafer-scale integration. In [A9], Zhang et al. introduce network-on-wafer interconnect that exploits photonic interconnect for wafer-scale systems and interconnects multi-chiplet graphics processing units together.

We hope this collection of articles provides a great overview on the challenges and opportunities from emerging system interconnect. The impact of communication and interconnect will only continue to grow as system integration and scalability increases, and we hope these articles encourage new research in this exciting area. We would like to thank all the authors who submitted articles to this Special Issue on Emerging System Interconnects and the anonymous reviewers, the Editor-in-Chief (Lizy Kurian John), and *IEEE Micro* staff (Jaclyn Martin) for their support in making this special issue happen.

APPENDIX: RELATED ARTICLES

[A1] D. D. Sharma, "Novel composable and scaleout architectures using Compute Express Link," *IEEE Micro*, vol. 43, no. 2, pp. 9–19, Mar./Apr. 2023, doi: 10.1109/MM.2023.3235972.

[A2] K. Kim et al., "SMT: Software-defined memory tiering for heterogeneous computing systems with CXL memory expander," *IEEE Micro*, vol. 43, no. 2, pp. 20–29, Mar./Apr. 2023, doi: 10.1109/MM.2023.3240774.

[A3] D. S. Berger et al., "Design tradeoffs in CXL-based memory pools for public cloud platforms," *IEEE Micro*, vol. 43, no. 2, pp. 30–38, Mar./Apr. 2023, doi: 10.1109/MM.2023.3241586.

[A4] M. Ha et al., "Dynamic capacity service for improving CXL pooled memory efficiency," *IEEE Micro*, vol. 43, no. 2, pp. 39–47, Mar./Apr. 2023, doi: 10.1109/MM.2023.3237756.

[A5] D. Gouk, M. Kwon, H. Bae, S. Lee, and M. Jung, "Memory pooling with CXL," *IEEE Micro*, vol. 43, no. 2, pp. 48–57, Mar./Apr. 2023, doi: 10.1109/MM.2023.3237491.

[A6] D. Boles, D. Waddington, and D. A. Roberts, "CXL-enabled enhanced memory functions," *IEEE Micro*, vol. 43, no. 2, pp. 58–65, Mar./Apr. 2023, doi: 10.1109/MM.2022.3229627.

[A7] M. Kwon, J. Jang, H. Choi, S. Lee, and M. Jung, "Failure tolerant training with persistent memory disaggregation over CXL," *IEEE Micro*, vol. 43, no. 2, pp. 66–75, Mar./Apr. 2023, doi: 10.1109/MM.2023.3237548.

[A8] D. D. Sharma, "System on a package innovations with Universal Chiplet Interconnect Express (UCIe) interconnect," *IEEE Micro*, vol. 43, no. 2, pp. 76–85, doi: 10.1109/MM.2023.3235770.

[A9] S. Zhang, "Photonic network-on-wafer for multi-chiplet GPUs," *IEEE Micro*, vol. 43, no. 2, pp. 86–95, Mar./Apr. 2023, doi: 10.1109/MM.2023.3237927.

JOHN KIM is a professor in the School of Electrical Engineering, Korea Advanced Institute of Science and Technology, Daejeon, 34141, South Korea. His research interests include computer architecture, interconnection networks, and large-scale systems. Kim received a Ph.D. degree from Stanford University, Stanford, CA, USA. He is a Senior Member of IEEE. Contact him at jjk12@kaist.edu.

NAM SUNG KIM is the W.J. "Jerry" Sanders III-Advanced Micro Devices Inc. Endowed Chair Professor at the University of Illinois, Urbana, IL, 61801, USA. His research interests include computer architecture and system software, focusing on memory, storage, and network subsystems. He is a Fellow of IEEE and ACM. Contact him at nskim@illinois.edu.