Sun XH, Li D. Special section on memory-centric system research for high-performance computing. JOURNAL OF COMPUTER SCIENCE AND TECHNOLOGY 36(1): 1–3 Jan. 2021. DOI 10.1007/s11390-021-0001-4

Preface

It is our great pleasure to announce the publication of this special section in Journal of Computer Science and Technology (JCST), Memory-Centric System Research for High-Performance Computing (HPC).

The growing disparity between CPU speed and memory speed, known as the memory-wall problem, has been a long-standing challenge in the computing industry. Several memory technologies and architectures including 3D-stacked memory, non-volatile random-access memory (NVRAM), memristor, hybrid software and hardware caches, etc. have been introduced in recent years to address the infamous memory-wall problem. This new expansion of memory technology provides more opportunities, but also further complicates the memory system design and optimization. Computer architecture and hardware system, operating systems, storage and file systems, programming stack, performance model and tools are being enhanced, augmented, and redesigned to utilize and support deeper and more diversified memory hierarchy systems for better performance and more. This is especially true for HPC data-intensive applications, where advanced technologies are always adopted first for best performances.

Memory systems in HPC are going through big changes recently, in order to accommodate ever-increasing datasets, reduce data movement, utilize advanced hardware devices, and improve application performance. Many research groups around the world are exploring emerging memory systems to enable efficient HPC systems and applications. For the HPC research community to build a general picture of the state-of-the-art and vibrant progress, we have taken an initiative to solicit contributions from some of the leading research groups in the field to form this special section. After two rounds of thorough review and revisions, with some withdrawn and declined, we end up with a special section including seven high-quality papers from Asia, Europe, and North America.

This special section begins with a review paper from an international team from China and Singapore. It presents "A Survey of Non-Volatile Main Memory Technologies: State-of-the-Arts, Practices, and Future Directions". This paper revisits the landscape of emerging non-volatile main memory (NVMM) technologies, and then surveys the state-of-the-art studies of NVMM technologies. It classifies those studies with a taxonomy according to different dimensions such as memory architectures, data persistence, performance improvement, energy saving, and wear leveling. To demonstrate the best practices in building NVMM systems, the paper introduces the authors' recent work of hybrid memory system designs from the dimensions of architectures, systems, and applications. The paper also presents the vision of future research directions of NVMMs and sheds some light on design challenges and opportunities.

There are three papers from China. The first paper, "Performance Evaluation of Memory-Centric ARMv8 Many-Core Architectures: A Case Study with Phytium 2000+", is authored by a team from the National University of Defense Technology. This paper presents a comprehensive performance evaluation of Phytium 2000+, an ARMv8-based 64-core architecture. It focuses on the cache and memory subsystems, analyzing the characteristics that impact HPC applications. Its evaluation results show that the ARMv8-based many-core system can deliver high performance for a wide range of scientific kernels. The second paper of the three, "A GPU-Accelerated In-Memory Metadata Management Scheme for Large-Scale Parallel File Systems" is authored by a team of researchers from the Sun Yat-sen University. This paper introduces a novel metadata server architecture, which offloads metadata tasks to GPU to improve the performance of metadata service. To fully utilize the GPU parallelism, a new in-memory data structure is introduced for the name space of file systems. The authors' experimental results demonstrate that the newly proposed GPU-based solution outperforms the conventional CPU-based scheme by

more than 50% under typical metadata operations. The third paper from China, "PIM-Align: A Processing-in-Memory Architecture for FM-Index Search Algorithm", is authored by a team from the Institute of Computing Technology, Chinese Academy of Sciences and University of Chinese Academy of Sciences. This paper introduces a processing-in-memory architecture to accelerate DNA alignment (the most critical and time-consuming step in genomic data analysis). To achieve memory-capacity proportional performance, the authors propose a lightweight message passing mechanism between memory vaults and a hardware prefetcher specialized for memory access patterns under their new processing-in-memory architecture. Evaluation results show that the proposed architecture can achieve 20x and 1820x speedup when compared with the best available ASIC solution and 32-thread CPU baseline, respectively.

Two papers are from the United States in this special section. The first paper, "A Study on Modeling and Optimization of Memory Systems", comes from a team in the Florida International University and Illinois Institute of Technology. Accesses Per Cycle (APC), Concurrent Average Memory Access Time (C-AMAT), and Layered PerformanceMatching (LPM) are three memory performance models that consider both data locality and memory assess concurrency. This paper reexamines the three models under one coherent mathematical framework. It divides the memory cycles at each memory layer into four distinct categories and uses them to recursively define the memory access latency and concurrency along the memory hierarchy. This new perspective offers new insights with a clear formulation of the memory performance. Consequently, the performance model can be easily understood and applied in engineering practices. The second paper, "Unimem: Runtime Data Management on Non-Volatile Memory-based Heterogeneous Main Memory for High Performance Computing", comes from the University of California Merced. This paper introduces a lightweight runtime solution that automatically and transparently manages data placement on heterogeneous memory systems without the requirement of any hardware modification and disruptive change to applications. By leveraging online profiling and performance models, the runtime solution characterizes memory access patterns associated with data objects and minimizes unnecessary data movement. The proposed solution effectively bridges the performance gap between NVM and DRAM. The authors demonstrate that using NVM to replace the majority of DRAM is a feasible solution for future HPC systems with the assistance of a software-based data management.

Finally, this special section presents one paper from Europe. The paper, "Usage Scenarios for Byte-Addressable Persistent Memory in High-Performance and Data Intensive Computing", comes from an international team at the University of Edinburgh and SVA Paderborn. This paper outlines some of the basic performance characteristics of Intel Optane DC persistent memory module (PMM), and explains how it can be configured and used to address the needs of memory and I/O intensive applications in HPC and data intensive domains. It shows that the flexibility of the Optane technology has the potential to be truly disruptive, not only because of the performance improvements it can deliver, but also because it allows systems to cater for wider range of applications on homogeneous hardware.

In summary, the memory systems in HPC are attracting more research efforts, from the perspectives of architecture, system software and applications. This special section offers a great glimpse of ongoing research efforts from various organizations in the world. We hope that readers find this special section inspiring and useful for their research and practice in HPC.

Acknowledgement

We would like to take the opportunity to express our great gratitude to all the authors for their contributions under the uneasy COVID-19 pandemic period, including those whose submissions were not accepted or withdrawn due to various reasons. We are also very thankful to all the reviewers for their valuable time and efforts. We are particularly grateful to the former Editor-in-Chief, Prof. Guo-Jie Li, and the incumbent Editor-in-Chief, Prof.

Zhi-Wei Xu, for hosting this special section. We also thank Ms. Feng-Di Shu and the other editorial office staff, for their tireless efforts and extreme patience in coordinating the communication and revisions throughout the entire review process.

Leading Editor

Xian-He Sun, Distinguished Professor, Illinois Institute of Technology, Chicago sun@iit.edu

Principal Guest Editor

Dong Li, Associate Professor, University of California, Merced dli35@ucmerced.edu

Xian-He Sun is a University Distinguished Professor of Computer Science at the Department of Computer Science in the Illinois Institute of Technology (IIT), Chicago. Before joining IIT, he worked at DoE Ames National Laboratory, at ICASE, NASA Langley Research Center, at Louisiana State University, Baton Rouge, and was an ASEE Fellow at Navy Research Laboratories. Dr. Sun is an IEEE Fellow and is known for his memory-bounded speedup model, also called Sun-Ni's Law, for scalable computing. His research interests include data-intensive high-performance computing, memory and I/O systems, software system for big data applications, and performance evaluation and optimization. He has over 250 publications and six patents in these areas. He is the Associate Chief Editor of IEEE Transactions on Parallel and Distributed Systems, a Golden Core member of the IEEE CS Society, and the past chair of the Computer Science Department at IIT. More information about Dr. Sun can be found at www.cs.iit.edu/∼sun/.

Dong Li is an associate professor at Department of Electrical Engineering and Computer Science (EECS) at University of California (UC), Merced. Previously, he was a research scientist at the Oak Ridge National Laboratory (ORNL), studying computer architecture and programming model for next generation supercomputer systems. Dong earned his Ph.D. degree in computer science from Virginia Tech. His research focuses on high-performance computing (HPC), and maintains a strong relevance to computer systems. The core theme of his research is to study how to enable scalable and efficient execution of enterprise and scientific applications on increasingly complex large-scale parallel systems. Dong received a CAREER Award from U.S. National Science Foundation in 2016, and an ORNL/CSMD Distinguished Contributor Award (2013). His paper in SC'14 was nominated as the Best Student Paper. He is also the lead PI for NVIDIA CUDA Research Center at UC Merced. He is a review board member of IEEE Transaction on Parallel and Distributed Systems (TPDS).