CCF Transactions on High Performance Computing (2020) 2:1–2
https://doi.org/10.1007/s42514-020-00026-9

EDITORIAL

Editorial for the special issue on storage system and technology

Dan Feng[1] · Hong Jiang[2] · Laurence T. Yang[3] · Xubin He[4] · André Brinkmann[5]

Published online: 17 March 2020

This issue focuses on the topic "Storage System and Technology". Data storage systems are an important part of high-performance computing (HPC). HPC cannot be separated from the support of high-performance storage systems and technologies. New storage technologies and techniques continue to be applied to HPC, such as non-volatile memory technologies, solid state storage, parallel I/O, storage performance and scalability, storage virtualization, and deduplication.

Five invited papers have been finally selected in this issue based on a peer-review procedure, which cover research progress on storage systems and technologies. There are two research papers on non-volatile memory (NVM) technologies. Emerging non-volatile memory has attractive characteristics such as DRAM-like low-latency together with the non-volatility of storage devices and is going on to solve the memory wall problem in HPC systems. The paper written by Chih Chieh Chou proposes vNVML, a user-space library that allows multiple applications to share and use a hybrid NVM-SSD-based persistent storage. The authors discuss the motivations and design details of the proposed vNVML, and evaluate its performance and overhead via realistic workloads and micro benchmark. Another paper written by Jiangkun Hu et al. covers consistency and efficiency of B + trees for NVM. The authors study and analyze the influence factors of designing B + trees on NVM with comprehensive evaluations and provide guidance on how to design efficient B + trees on NVM. Redundant arrays of independent disks (RAID) are the most wide-spread technology to secure data against failures in data centers. The rapid increase of data, however, increases the pressure on RAID scalability. Chentao Wu proposes an approximate intelligent redistribution approach, which predicts dynamic data access patterns from running workloads and minimizes the movement of data. Their experimental results show that their approach reduces data migration and speeds up the scaling process. It is necessary to reevaluate all levels of the storage hierarchy to optimize software performance when emerging data storage devices are adopted in a supercomputing center. The paper written by Alessandro Elias proposes a simulation technique to emulate storage devices, called Freezing Time, which pauses a virtual machine to manipulate its clock and hide the real I/O completion time. The evaluation shows that this technique is able to emulate disks with RAM-like speeds. The last paper written by Yutong Lu analyzes the problems of performance degradation and scalability for data movements caused by I/O stack virtualization, and proposes concrete methods to improve the performance of HPC collective data movements and bursty asynchronous data movements by dynamically using the modest mode.

Finally, we would like to thank to all reviewers and all the authors for their contributions. It has been the tight cooperation between them that makes this issue ready for readers. We look forward to more and more papers of high quality appearing in the CCF THPC journal.

✉ Dan Feng
dfeng@hust.edu.cn

Hong Jiang
hong.jiang@uta.edu

Laurence T. Yang
ltyang@stfx.ca

Xubin He
xubin.he@temple.edu

[1] Huazhong University of Science and Technology, Wuhan 430074, China

[2] University of Texas at Arlington, Arlington, USA

[3] St Francis Xavier University, Antigonish, Canada

[4] Temple University, Philadelphia, USA

[5] Johannes Gutenberg-Universität Mainz, Mainz, Germany

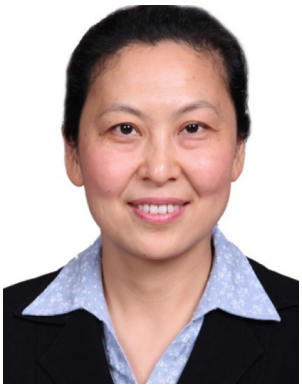

Dan Feng received the B.E., M.E., and Ph.D. degrees in computer science and technology from Huazhong University of Science and Technology (HUST), Wuhan, China, in 1991, 1994, and 1997, respectively. She is a Professor and the Dean of the School of Computer Science and Technology, HUST. Her research interests include computer architecture, Non-Volatile memory technology, distributed and parallel file system, and massive storage system. She has more than 100 publications in major journals and international conferences, including IEEE TC, IEEE TPDS, ACM-TOS, FAST, USENIX ATC, EuroSys, ICDCS, HPDC, SC, ICS, IPDPS, DAC and DATE. Dr. Feng has served as the program committees of multiple international conferences, including SC 2011, 2013 and MSST 2012, 2015.

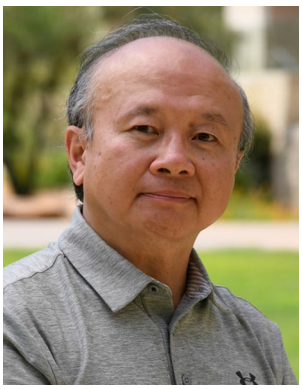

Hong Jiang received the B.Sc. degree in Computer Engineering from Huazhong University of Science and Technology, Wuhan, China; the M.A.Sc. degree in Computer Engineering from the University of Toronto, Toronto, Canada; and the PhD degree in Computer Science from the Texas A&M University, College Station, Texas, USA. He is currently Chair and Wendell H. Nedderman Endowed Professor of Computer Science and Engineering Department at the University of Texas at Arlington. Prior to joining UTA, he served as a Program Director at National Science Foundation (2013.1–2015.8) and he was at University of Nebraska-Lincoln since 1991, where he was Willa Cather Professor of Computer Science and Engineering. His present research interests include computer architecture, computer storage systems and parallel I/O, high-performance computing, big data computing, and cloud and edge computing. He has over 300 publications in major journals and international Conferences in these areas. He is a Fellow of IEEE, and Member of ACM.

Laurence T. Yang is a professor and the W.F. James Research Chair in Department of Computer Science, St. Francis Xavier University, Canada and a professor in School of Computer Science and Technology, Huazhong University of Science and Technology, China. He received the B.E. degree in computer science and technology and B.Sc. degree in applied physics both from Tsinghua University, China, and the Ph.D degree in computer science from University of Victoria, Canada. His research interests include parallel, distributed and cloud computing, embedded and ubiquitous/pervasive computing, big data and cyber-physical-social systems. He is a Clarivate Analytics Highly Cited Researcher, an elected Fellow of IEEE, the Canadian Academy of Engineering and the Engineering Institute of Canada.

Xubin He received the BS and MS degrees in computer science from Huazhong University of Science and Technology, China, in 1995 and 1997, respectively, and the PhD degree in electrical engineering from the University of Rhode Island, Kingston, Rhode Island, in 2002. He is currently a professor in the Department of Computer and Information Sciences, Temple University, Philadelphia, Pennsylvania. His research interests include computer architecture, data storage systems, virtualization, and high availability computing. He received the Ralph E. Power Junior Faculty Enhancement Award in 2004 and the Sigma Xi Research Award (TTU Chapter) in 2005 and 2010. He is a senior member of the IEEE, a member of the IEEE Computer Society and USENIX.

André Brinkmann is a full professor at the Computer Science Department of JGU and head of the ZDV (since 2011). He received his Ph.D. in electrical engineering in 2004 from the Paderborn University and was an assistant professor in the computer science department of the Paderborn University from 2008 to 2011. Furthermore, he has been the managing director of the Paderborn Centre for Parallel Computing PC^2 during this time frame. His research interests focus on the application of algorithm engineering techniques in the area of data centre management, cloud computing, and storage systems. He has published more than 100 papers in renowned conferences and journals and is an associated editor of the ACM Transactions on Storage as well as a steering committee member of the IEEE International Conference on Networking, Architecture, and Storage (NAS). He is a member of the advisory board of the French Grid'5000.