



# Editorial for the special issue on high performance distributed computing

Minyi Guo<sup>1</sup> · Guihai Chen<sup>1</sup> · Xiaofei Liao<sup>2</sup> · Long Zheng<sup>2</sup>

Published online: 8 April 2021  
© China Computer Federation (CCF) 2021

With the advent of the Big Data era, the demand for storing, processing, and analyzing this vast and growing amount of data has emerged in the market. Since a single node cannot cope with its complexity requirements, high-performance systems typically operate in a distributed environment. Today, high-performance distributed computing systems are the foundations of many important infrastructures. However, the diverse development of hardware platforms, the spurt of data growth, and the rapid changes in applications are increasingly challenging resource management, energy efficiency, performance tuning, scalability, and fault tolerance.

We have five invited papers selected for this special issue based on a peer-review procedure, which covers a few different aspects related to the memory system and deep neural network. Specifically, the first two papers discuss the synchronization of memory copy mechanism and configuration optimization of distributed storage systems, respectively. The next two papers focus on the performance and programmability of distributed neural network systems. In addition, there is a paper exploiting transfer learning for pneumoconiosis detection, which is an interesting and valuable attempt.

- The paper written by Zhenke Chen et al. presents an asynchronous memory copying mechanism based on the Intel I/OAT engine to reduce the synchronization overheads. It guarantees the consistency of I/OAT queues,

achieves scalability for buffering copying requests to I/OAT channels concurrently, and gets parallelism for parallel execution of multiple I/OAT channels. Experiments are conducted to verify its performance superiority to the existing copying mechanism for different copy size in NUMA scenarios.

- The paper written by Xin Li et al. analyzes the data recovery mechanism in HDFS and finds that the configuration of replication tasks can affect the data recovery significantly. They propose a strategy to dynamically optimize replicators and resource parameters based on the existing data recovery strategy. The results show that it can effectively speed up the recovery time by 71%. It also improves the execution performance of MapReduce jobs from 28 to 59% through optimized replication factors and resource parameters.
- The paper written by Haoyu Cai et al. designs a complete solution to the Kubeflow pipeline problem. It verifies the feasibility of the solution through an example so that Kubernetes no longer depends on Google Cloud Service and better supports machine learning applications based on Pytorch. Based on Kubeflow, a containerized distributed training program is implemented, making containerized distributed deployment easier. Experimental results with the deployment of yolov3 inferred services illustrate the containerization and microservices with promising performance.
- The paper written by Qiang Qi et al. proposes a network bandwidth resource allocation strategy for distributed deep neural network (DNN) training tasks to mitigate the network resource contention and alleviate the performance variation of training jobs. It first acquires the communication data information, including the layer index and its class based on job profiling, and then calculates the weights of two co-located tasks according to the class information of communication data transferred by parameter server and worker tasks. The experimental

---

✉ Long Zheng  
longzh@hust.edu.cn

Minyi Guo  
guo-my@cs.sjtu.edu.cn

Guihai Chen  
gchen@cs.sjtu.edu.cn

Xiaofei Liao  
xfliao@hust.edu.cn

<sup>1</sup> Shanghai Jiao Tong University, Shanghai, China

<sup>2</sup> Huazhong University of Science and Technology, Wuhan, China

results show its performance superiority to existing systems.

- The paper written by Ran Zheng et al. studies pneumoconiosis identification from the perspective of CNN-based transfer learning, which originally required an expert doctor for a diagnosis. They propose a transfer learning pattern to solve the problem of scarce annotated pneumoconiosis data. Some preprocessing methods are presented to improve X-ray films' qualities, including image denoising, lung segmentation, and data amplification. The experimental results demonstrate the effectiveness of learning in the identification of pneumoconiosis, even with limited training data.

We would like to take this chance to thank all the authors and the reviewers for their brilliant contribution to this special issue of CCF THPC. With their great efforts, we can put together the five research papers that discuss different topics and present different ideas that help promote high-performance distributed computing.



**Minyi Guo** received the Ph.D degree in computer science from the University of Tsukuba, Japan. He is currently Zhiyuan Chair professor and the head of the Department of Computer Science and Engineering, Shanghai Jiao Tong University, China. His present research interests include parallel/distributed computing, compiler optimizations, embedded systems, pervasive computing, big data, and cloud computing. He is currently on the editorial board of IEEE Transactions on Parallel and Distributed

Systems, IEEE Transactions on Cloud Computing, and Journal of Parallel and Distributed Computing. He is a fellow of CCF and IEEE.



**Guihai Chen** received the B.S. degree from Nanjing University, in 1984, the ME degree from Southeast University, in 1987, and the Ph.D degree from the University of Hong Kong, in 1997. He is a distinguished professor of Shanghai Jiao Tong University, China. He had been invited as a visiting professor by many universities including the Kyushu Institute of Technology, Japan, in 1998, the University of Queensland, Australia, in 2000,

and Wayne State University during September 2001 to August 2003. He has a wide range of research interests with focus on sensor network, peer-to-peer computing, high-performance computer architecture, and combinatorics. He has published more than 200 peer-reviewed papers, and more than 120 of them are in well-archived international journals such as the IEEE Transactions on Parallel and Distributed Systems, the Journal of Parallel and Distributed Computing, the Wireless Network, The Computer Journal, the International Journal of Foundations of Computer Science, and Performance Evaluation, and also in well-known conference proceedings such as HPCA, MOBIHOC, INFOCOM, ICNP, ICPP, IPDPS, and ICDCS. He is a senior member of the IEEE.



**Xiaofei Liao** received his Ph.D. degree in computer science and engineering from Huazhong University of Science and Technology (HUST), Wuhan, in 2005. He is currently a professor in the School of Computer Science and Technology at HUST, Wuhan. He has served as a reviewer for many conferences and journal papers. His research interests are in the areas of system software, P2P system, cluster computing and streaming services. He is a member of IEEE and the IEEE Computer

Society.



**Long Zheng** is now an associate professor in the School of Computer Science and Technology, Huazhong University of Science and Technology (HUST), Wuhan. He received his Ph.D. degree in computer engineering at HUST, Wuhan, in 2016. His current research interests include configurable computer architecture and runtime systems.