

Data Storage and Management (DSM)

Time: Monday-Tuesday, December 19-20, 2022

Online Conference Room: Zoom Code: 937 6341 9074 Password: 535129

Chair

Hong JIANG

The University of Texas at Arlington

Co-chairs

Guangyan ZHANG

Tsinghua University

Yu HUA

Huazhong University of Science and Technology

Hua WANG

Huazhong University of Science and Technology

Program Committee Members

Helei CUI, Northwestern Polytechnical University

Dan FENG, Huazhong University of Science and Technology

Dejun JIANG, Institute of Computing Technology, CAS

Duo LIU, Chongqing University

Zili SHAO, The Chinese University of Hong Kong

Jiwu SHU, Tsinghua University

Yang WANG, Ohio State University

Yi WANG, Shenzhen University

Nong XIAO, Sun Yat-sen University

Changsheng XIE, Huazhong University of Science and Technology

Chun XUE, City University of Hong Kong

Ke ZHOU, Huazhong University of Science and Technology

14:00-15:20 December 19**Presider: Guangyan ZHANG, Tsinghua University**

Time	Speaker	Talk Title
14:00-14:40	Tei-Wei KUO (Keynote)	Rethinking of computing - Memory-centric computing
14:40-15:20	Andre Brinkmann (Keynote)	Improving HPC I/O by coupling Lustre with ad hoc file systems

15:35-17:10 December 19**Presider: Guangyan ZHANG, Tsinghua University**

15:35-16:00	Zili SHAO (Invited)	Memory/Storage optimization for small/big systems
16:00-16:25	Chun XUE (Invited)	On optimizing mobile memory management

Best Student Paper

16:25-16:40	Yaodong HAN	Research on workload prediction in cloud block storage system
16:40-16:55	Hui CHEN	A new fast testing method for actuator characteristics of optical disk drives
16:55-17:10	Chenwei LIU	Task scheduling strategy based on energy consumption optimization in cloud data center

08:30-10:15 December 20**Presider: Hong JIANG, University of Texas at Arlington**

Time	Speaker	Talk Title
08:30-09:10	David Hung-Chang DU (Keynote)	Computing in-X for supporting big data applications
09:10-09:50	Calvin LIN (Keynote)	Prediction mechanisms in computer architecture

09:50-10:15	Yang WANG (Invited)	TBD
10:30-11:45 December 20		
Presider:		
10:30-10:55	Chenlin MA (Invited)	TBD
10:55-11:20	Yujuan TAN (Invited)	SAPredictor: A simple and accurate self-adaptive predictor for hierarchical hybrid memory system
11:20-11:45	Helei CUI (Invited)	Secure deduplication over encrypted data: New perspectives and directions
12:00-13:30	Lunch Time	
14:00-15:45 December 20		
Presider: Yu HUA, Huazhong University of Science and Technology		
Time	Speaker	Talk Title
14:00-14:25	Dejun JIANG (Invited)	NapFS: A high-performance NUMA-aware PM file system
14:25-14:50	Xing XU (Invited)	Recent advances in cross-modal video event retrieval and localization
14:50-15:15	Ji ZHANG (Invited)	TBD
15:15-15:30	Yibin GU	LPCA: Learned MRC profiling based cache allocation for file storage systems
15:30-15:45	Chen XIE	Design of optical disk drive servo channel

Monday, December 19, 2022(14:00-15:20)

Presider: Guangyan ZHANG, Tsinghua University

14:00-14:40 (Keynote)

Rethinking of computing - Memory-centric computing

Tei-Wei KUO

Taiwan University

Abstract: Flash memory opens a window of opportunities to a new world of computing over 20 years ago. Since then, storage devices gain their momentum in performance, energy, and even access behaviors. With over 1000 times in performance improvement over storage in recent years, there is another wave of adventure in removing traditional I/O bottlenecks in computer designs. In this talk, I shall address the opportunities of new system architectures in computing. In particular, hybrid modules of DRAM and non-volatile memory (NVM) and all NVM-based main memory will be considered. I would also comment on a joint management framework of host/CPU and a hybrid memory module to break down the great memory wall by bridging the process information gap between host/CPU and a hybrid memory module. Before the conclusion of this talk, I will further extend the idea of memory-centric computing to neuromorphic computing which empower memory chips to own new capabilities in computing. In particular, I shall address challenges in in-memory computing in realizing deep learning.

14:40-15:20 (Keynote)

Improving HPC I/O by coupling Lustre with ad hoc file systems

Andre Brinkmann

Johannes Gutenberg-University Mainz

Abstract: Today's high-performance computing (HPC) centers use parallel file systems (PFSs) to satisfy the I/O needs of a large variety of applications. The underlying backend storage is often based on slow magnetic disks, while HPC compute nodes are equipped with fast storage devices like SSDs or non-volatile main memory (NVMM). Parallel file systems can leverage this fast storage by aggregating their performance and capacity to build a caching layer.

Lustre, as an example for a wide-spread PFS, offers persistent client caching for Lustre (LPCC) on top of its hierarchical storage management to build such a caching layer. LPCC offers two cache modes: read-only and read-write. However, both are restricted to being used only as local caches and clients still constantly communicate over the network with a remote metadata server to ensure a consistent namespace.

In this talk, we will introduce LPCC and discuss how ad hoc file systems can help to provide a distributed cache building on LPCC's existing functionality. We will propose a new architecture that couples Lustre and our ad-hoc file system GekkoFS. Coupling Lustre with GekkoFS, we will offer both concurrent I/O support from multiple nodes for LPCC and accelerated metadata performance while still maintaining namespace consistency. Furthermore, GekkoFS can adapt to the specific I/O requirements of applications and reduce cross-application interference.

Tuesday, December 20, 2022(15:35-17:10)

Presider: Guangyan ZHANG, Tsinghua University

15:35-16:00 (Invited)

Memory/Storage optimization for small/big systems

Zili SHAO

The Chinese University of Hong Kong

Abstract: Memory/storage optimization is one of the most critical issues in computer systems. In this talk, I will first summarize our work in optimizing memory/storage systems for embedded and big data applications. Then, I will present our recent work in optimizing LSM-tree based Relational Databases.

16:00-16:25 (Invited)

On optimizing mobile memory management

Chun XUE

City University of Hong Kong

Abstract: Current Mobile Operating System such as Android inherits Linux kernel. As a result, designs that were targeting servers are now applied in mobile devices. In this series of work, through analysing mobile APP characteristics on files, memory, and storage usage, we found that APPs have its own unique characteristics which differ from applications on servers. These differences present new optimization opportunities in mobile memory management. In this talk, we will present several works that improve user experiences for mobile devices based on APP characterization.

16:25-16:40 (Oral)

Research on workload prediction in cloud block storage system

Yaodong HAN

Huazhong University of Science and Technology

Abstract: This paper proposes a cloud disk load prediction method, which can improve resource utilization by 7.5~20% compared with the previous method

16:40-16:55 (Oral)

A New Fast Testing Method for Actuator Characteristics of Optical Disk Drives

Hui CHEN

Huazhong University of Science and Technology

Abstract: Optical disk drive is a way to realize optical storage.

16:55-17:10 (Oral)

Task scheduling strategy based on energy consumption optimization in cloud data center

Chen-wei LIU

North Minzu University

Abstract: The energy consumption problem has attracted the attention of academia. A chaotic mapping adaptive particle swarm optimization algorithm based on opposition-based learning is proposed. Balance energy consumption and user costs by shortening task completion time.

Tuesday, December 20, 2022(08:30-10:15)

Presider: Hong JIANG, University of Texas at Arlington

08:30-09:10 (Keynote)

Computing in-X for supporting big data applications

David Hung-Chang DU

University of Minnesota

Abstract: Many challenges exist in big data era. We intend to fully utilize the available data for making critical decisions. In this talk, we introduce several potential improvements for meeting these big data challenges, including how to use non-volatile memory, active storage devices, and software-defined networks, in our current computing environment. We will especially introduce computing in-storage and in-networking in addition to computing in-memory.

09:10-09:50 (Keynote)

Prediction mechanisms in computer architecture

Calvin LIN

University of Texas at Austin

TBD

09:50-10:15 (Invited)

Talk Title

Yang WANG

Ohio State University

TBD

Tuesday, December 20, 2022(10:30-11:45)

Presider:

10:30-10:55 (Invited)

Talk Title

Chenlin MA

Shenzhen University

TBD

10:55-11:20 (Invited)**TBD**

Yujuan TAN

Chongqing University

Abstract: In a hybrid memory system using DRAM as the NVM cache, DRAM and NVM can be accessed in serial or parallel mode. However, using either mode alone will bring access latency and bandwidth problems. To solve this problem, we integrate these two access modes and design a simple but accurate predictor (called SAPredictor) to help choose the appropriate access mode, thereby avoiding long access latency and bandwidth problems to improve memory performance.

11:20-11:45 (Invited)**Secure deduplication over encrypted data: New perspectives and directions**

Helei CUI

Northwestern Polytechnical University

Abstract: In the past decades, secure deduplication over encrypted data has been studied intensively for achieving high storage utilization and strong privacy protection simultaneously. Now, we'd like to explore how to apply such techniques in other different contexts, explicitly for accelerating SGX-enabled applications and improving decentralized cloud storage networks.

Tuesday, December 20, 2022(14:00-15:45)**President: Yu HUA, Huazhong University of Science and Technology****14:00-14:25 (Invited)****NapFS: A high-performance NUMA-aware PM file system**

Dejun JIANG

Institute of Computing Technology, CAS

Abstract: In this talk, I will present NapFS, a high-performance NUMA-aware persistent memory file system. I first show NUMA impacts and design principles when building PM file systems. Then, I introduce the key techniques adopted in NapFS, including IO delegation and per-socket local FS. I finally show the efficiency of NapFS.

14:25-14:50 (Invited)**Recent advances in cross-modal video event retrieval and localization**

Xing XU

University of Electronic Science and Technology of China

Abstract: We dedicate to understanding and retrieving desired content/event from complicated videos, improving users' experience in the current multimedia information retrieval systems. In this talk, I will introduce several advanced techniques in multimodal representation, fusion, alignment, and collaborative learning in practical cross-modal video event retrieval and localization.

14:50-15:15 (Invited)

TBD

Ji ZHANG

Huawei Technologies Co., Ltd.

TBD

15:15-15:30 (Oral)

LPCA: Learned MRC profiling based cache allocation for file storage systems

Yibin GU

Huazhong University of Science and Technology

Abstract: This paper proposed a learned MRC profiling-based cache allocation scheme named LPCA, taking the MRC as guides to allocate cache space for multiple data streams, thus improving the overall performance.

15:30-15:45 (Oral)

Design of optical disk drive servo channel

Chen XIE

Huazhong University of Science and Technology

TBD