CCF Transactions on High Performance Computing (2022) 4:231–232
https://doi.org/10.1007/s42514-022-00122-y

EDITORIAL

Editorial for the special issue on intelligent storage and edge computing

Jiwu Shu[1] · Ke Zhou[2] · Xu Chen[3]

Published online: 29 September 2022

With the rapid development and accelerating convergence of technical sectors such as internet-of-things (IoT) and 5G networks, the recent years have witnessed a massive amount of data being generated all the time. The profound value of these big data can be fully unleashed only if they can be accessed, organized and analyzed efficiently. As a result, intelligent storage and edge computing are emerging as novel enabling technologies, which aim at managing and serving data in an intelligent manner and supporting real-time data processing life-cycle efficiently. Intelligent storage and edge computing can be deployed in a multitude of diverse environments and enable flexible data moving and processing over the continuum from devices on-premises, edge servers to cloud datacenters, and can also well adapt to complex system changes. This issue focuses on the novel ideas, methods, designs or experimental implementations and evaluations of intelligent storage and edge computing.

We have seven invited papers selected for this special issue based on a peer-review procedure, which cover a variety of aspects that relate to intelligent storage and edge computing mentioned above.

The first part of the special issue focuses on the theme of intelligent storage, which presents one survey on AI for storage and four recent intelligent storage based achievements including three in SSD and one in cloud storage.

✉ Xu Chen
chenxu35@mail.sysu.edu.cn

Jiwu Shu
shujw@tsinghua.edu.cn

Ke Zhou
zhke@hust.edu.cn

[1] Tsinghua University/Xiamen University, Beijing, China

[2] Huazhong University of Science and Technology, Hongshan, China

[3] Sun Yat-sen University, Guangzhou, China

- The paper written by Prof. Hua Wang and her team presents a comprehensive survey on "AI for Storage", and categorizes storage research employing intelligent algorithms into Architecture-oriented, Data-specific, and Operation & Maintenance according to public literatures in recent years. Based on this classification, they fine-categorize all of the studies by application scenarios and elaborate on their development history in order to provide guidelines for future research on how to employ AI technologies to storage optimization. This paper is the first survey on "AI for storage" as we known.
- The paper written by Prof. Shi Liang and his team proposes a read latency variation aware performance optimization scheme, RLV, to accelerate both data and metadata accesses to maximize the read performance and reduce the tail latency. Experimental results show that the proposed method can improve the read performance by 45.7\% on average compared with state-of-the-art works and significantly reduce the tail latency at 95-99.99 percentiles.
- The paper written by Prof. Qiao Li and his team proposes to minimize the number of refresh operations by exploiting the optimal RBER (ORBER). They develop an ORBER model by conducting evaluations and analyses on a set of real 3D NAND flash chips. Experiments show that, within the lifetime of 3D NAND flash memory, the proposed method can averagely reduce 75% of the P/E cycles consumed by refresh operations and improve lifetime by 2.5Xwith marginal overhead, compared to the traditional refresh scheme.
- The paper written by Prof. Huawei Li and her team proposed a Cognitive SSD + system and used it to build a deep learning-based unstructured data retrieval engine. In the proposed cognitive SSD+, a flash-accessing accelerator, DHS-x, is placed by the side of flash devices to achieve near-data deep learning and hybrid data search (DHS). Experimental results on the FPGA-based prototype reveal that the cognitive SSD + running

the proposed DHS framework achieves performance speedup of 3.48X in comparison with the counterpart based on conventional CPU and storage system, and it reduces the overall system energy consumption by up to 4.89X and 1.77X respectively when compared to CPU and GPU based solutions that deliver comparable performance.

- The paper written by Prof. Zhiguang Chen and his team proposes PrecisePeriod, a precise periodicity detection algorithm customized for multi-tenant cloud storage. It removes outliers through data preprocessing, employs the discrete wavelet transform to remove high-frequency noise while keeping frequency domain information, computes the candidate periodicity queue using the autocorrelation function, and determines precise period through periodicity verification. The evaluation shows that the PrecisePeriod scheduling significantly reduces tail latency while only bringing 1 -2% overhead.

The second part of the special issue focuses on the theme of edge computing, which presents the recent achievements in edge computing based connected vehicle system and content popularity prediction based edge caching.

- The paper written by Prof. Weisong Shi and his team targets on leveraging edge computing for enabling smart connected and autonomous vehicle (CAV) services. An extensible and flexible middleware of EdgeWare is developed to manage the execution of vehicle services, which can greatly accelerate the execution of services and save CPU and memory utilization.
- The paper written by Prof. Xiaofei Wang and his team falls into the hot topic of edge content caching. An edge cache node architecture is introduced, and accordingly a novel cache replacement strategy, named Cluster-Based Content Caching (CBCC) is developed, which can efficiently extract the content feature to predict popularity for cache replacement, reduce the pressure on the backbone and further improve user satisfaction.

We would like to take this chance to thank all the authors and the reviewers for their splendid contribution to this special issue of CCF THPC.

Jiwu Shu is a Professor of Tsinghua University, China. His current research interests include intelligent storage system, edge storage systems, network storage systems, non-volatile memory, processing in memory, storage security and reliability. He has published more than 100 relevant academic papers, including papers from international conferences such as SOSP, OSDI, SIGMOD, FAST, USENIX ATC, ISCA, ASPLOS, MICRO, SC, VLDB and ICDE, as well as papers from international journals such as IEEE/ACM Trans. He currently serves as associate editors for IEEE Transactions on Computers and ACM Transactions on Storage.

Ke Zhou is a Professor of Huazhong University of Science and Technology, China. His current research interests include AI for data storage and management, cloud storage, storage security. He has published more than 100 relevant academic papers, including papers from international conferences such as SIGMOD, FAST, USENIX ATC, DAC, MM, ICPP and MSST, as well as papers from international journals such as IEEE/ACM Trans. He currently serves as the director of intelligent cloud storage joint laboratory of HUST and Tencent.

Xu Chen is a Professor of Sun Yat-sen University, China. He received the IEEE Computer Society Best Paper Awards Runner-up, IEEE Communication Society Asia-Pacific Outstanding Young Researcher Award, IEEE INFOCOM Best Paper Runner-up Award and IEEE ICC Best Paper Award. He currently serves as associate editors for IEEE Transactions on Wireless Communications, IEEE Transactions on Vehicular Technology and IEEE Internet of Things Journal.