

Project on Tiered Storage Architecture Based on New Storage Technologies

Mingzhe Du

Project summary

1. Overview

The rapidly increasing data drives the development of new storage technologies like Solid State Disk, PCM, STT-RAM, Memristor, DWM, etc. Those kind of new storage devices provide features such as processing power, non-volatile storage. And the boundary between persistent storage, volatile storage and processing unit is increasingly getting blurred. New storage technologies are of paramount importance in improving the performance and capacity of datacenter. However, the heterogeneity of storage devices brings great challenges to the design of storage system architecture. How to make full use of the unique characteristics of new memory devices and maximize their efficacy is a key technical issue that needs to be solved, and it is also a hot research issue. This project is driven by this goal and focuses on the key technologies of data center storage system.

2. Broader Impact

Compared with traditional data centers, in the era of big data, data center storage systems have different requirements for networks such as delay and performance. These demands have fundamentally changed the network computing model from computation-centric one to datacenter-centric one. Correspondingly, storage technology has also undergone tremendous development with the emergence of block storage, file storage, and object storage. However, in this evolutionary development process, how to efficiently support data storage, management and processing, and shorten the distance from computation to data is always most important for quality of service.

Project description

1. Introduction

In recent years, new non-volatile memory (NVM) technologies have developed rapidly, such as phase change memory (PCM), spin transfer torque memory (STT-RAM), memristor (Memristor), track memory (DWM), etc. In addition to the same read and write performance compared to DRAM, these new NVM memories also provide a byte-level addressing feature, or memory-level access, that flash does not have, thus subversively changing the external memory. External byte-addressable storage not only changes the IO interface, but also changes the way to manage the storage and file system. The block-oriented optimization technology in traditional file systems is likely to be no longer effective. Storage-centric technologies cover a wider range of new memory devices, from block-addressable flash SSDs to byte-addressable NVM devices, all of which can be used as block devices. However, due to the limitation in cost, life expectation, I/O characteristics of new storage media, it is unlikely to completely replace

traditional disks by new ones. In fact, it is not necessary to replace traditional disk due to the diversity of data access features. Therefore, in current system architecture, traditional and new storage media are mixed together. Now the research about storage architecture mainly focuses on HDD and DRAM, which is not applicable when considering the new storage devices. For example, storage-centric technologies, i.e., memory-class storage can be integrated into block-based storage system by a drop-in way without changing existing architecture. While memory-centric technology, i.e., storage-class memory requires new programming model and system to take full advantage of byte-addressing. This project aims to solve the performance problem in the heterogeneous storage system, and to accelerate the upgrade of data centers. The main goals of this project are:

- (1) Research on memory-level storage architecture based on new storage technology;
- (2) Typical algorithms and theoretical studies of in-storage computations;
- (3) Research on distributed storage networking technology under memory-level storage architecture;
- (4) Efficient file system based on the new heterogeneous storage architecture.

References:

- [1] M.K Qureshi, M.M Franceschini, L.A Lastras-Montano, "Improving read performance of Phase Change Memories via Write Cancellation and Write Pausing", Proceedings of the IEEE 16th International Symposium on High Performance Computer Architecture (HPCA). Bangalore, India, 2010.
- [2] Joo Y, Niu D, Dong X, et al. "Energy-and endurance-aware design of phase change memory caches" Proceedings of the Conference on Design, Automation and Testin Europe. Leuven, Belgium, 2010: 136-141.
- [3] Yun J, Lee S. Yoo S. "Bloom filter-based dynamic wear leveling for phase-change RAM", Proceedings of the 42nd Annual IEEE/ACM International Symposium on Microarchitecture. New York, USA, 2009: 347-357.
- [4] M.K Qureshi, M.M Franceschini, L.A Lastras-Montano, et al. "Morphable memory system:A robust architecture for exploiting multi-level phase change memories", Proceedings of the 37th Annual International Symposium on Computer Architecture. New York, USA, 2010:153-162.
- [5] Lam CH. "Storage class memory", Proceedings of the 10th IEEE International Conf. on Solid-State and Integrated Circuit Technology (ICSICT). Shanghai, China, 2010, pp:1080-1083.
- [6] A. Caulfield, L. Grupp, and S. Swanson, "Gordon: Using Flash Memory to Build Fast, Power-efficient Clusters for Data-intensive Applications," in

Proceedings of the Fourteenth International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '09), 2009.

[7] J. Ouyang, S. Lin, S. Jiang, Z. Hou, Y. Wang, and Y. Wang, "SDF: Software-Defined Flash for Web-Scale Internet Storage Systems," in Proceedings of the 19th International Conference on Architectural Support for Programming Languages and Operating Systems, 2014, pp. 471-484.

[8] J. Ou, J. Shu, Y. Lu, L. Yi, and W. Wang, "EDM: An Endurance-Aware Data Migration

Scheme for Load Balancing in SSD Storage Clusters," in Proceedings of 28th IEEE International Parallel and Distributed Processing Symposium, 2014.

[9] Y. Lu, J. Shu, and W. Wang, "ReconFS: A Reconstructable File System on Flash Storage," in Proceedings of the 12th USENIX Conference on File and Storage Technologies (FAST '14), 2014, pp. 75-88.

[10] P. Wang, G. Sun, S. Jiang, J. Ouyang, S. Lin, C. Zhang, et al., "An Efficient Design and Implementation of LSM-Tree Based Key-Value Store on Open-Channel SSD," in Proceedings of the Ninth European Conference on Computer Systems, 2014, pp. 1-14.