

Report on the International Workshop on Big Data Management on Emerging Hardware (HardBD 2015)

Shimin Chen
Institute of Computing Technology
Chinese Academy of Sciences
chensm@ict.ac.cn

Qiong Luo
Hong Kong University of
Science and Technology
luo@cse.ust.hk

Xiaofeng Meng
School of Information
Renmin University of China
mengxf@ruc.edu.cn

1. INTRODUCTION

HardBD 2015 (International Workshop on Big Data Management on Emerging Hardware) was held in Seoul, Korea on April 13, 2015, in conjunction with ICDE 2015 (the 31st IEEE International Conference on Data Engineering) [1]. The aim of this half-day workshop is to bring together researchers, practitioners, system administrators, and others interested in management of big data over new hardware platforms to share their perspectives, and to discuss and identify future directions and challenges in this area.

Data properties and hardware characteristics are two key aspects for efficient data management. A clear trend in the first aspect, data properties, is the increasing demand to manage and process *Big Data* in both enterprise and consumer applications, characterized by the fast evolution of *Big Data Systems*. Examples of big data systems include NoSQL storage systems, MapReduce/Hadoop, data analytics platforms, search and indexing platforms, messaging infrastructures, event log processing systems, as well as novel extensions to relational database systems. These systems address needs for processing structured, semi-structured, and unstructured data across a wide spectrum of domains such as web, social networks, enterprise, mobile computing, sensor networks, multimedia/streaming, cyber-physical and high performance systems, and for a great many application areas such as e-commerce, finance, health care, transportation, telecommunication, and scientific computing. At the same time, the second aspect, hardware characteristics, is undergoing rapid changes, imposing new challenges for the efficient utilization of hardware resources. Recent trends include massive multi-core processing systems, high performance co-processors, very large main memory, emerging non-volatile memory technology, fast networking components, big computing clusters, and large data centers that consume massive amount of energy. Utilizing new hardware technologies for efficient Big Data management is of urgent importance.

The program committee accepted four regular papers that cover a variety of interesting topics, by authors from

China, Germany, Korea, and USA. The program also features a keynote speech by Sangyeun Cho, VP at Samsung for advanced solutions research and development, on recent advances in Flash solutions.

In the following, we will overview the keynote talk in Section 2 and the research papers in Section 3, and summarize the workshop in Section 4.

2. KEYNOTE TALK

The first session is keynote talk. Sangyeun Cho, our keynote speaker, became a tenured associate professor at the University of Pittsburgh in 2010. He recently joined Samsung's Memory Division as VP for advanced solutions research and development. His research interests are in the area of computer architecture and systems with particular focus on performance, power and reliability of memory and storage hierarchy design for next-generation data centers.

The speaker first described recent trends of Flash technology. Flash encounters significant challenges in scaling down its feature sizes. As density increases, cell interference goes up and the chance of data corruption increases. To address this problem (if temporarily), the memory industry has developed 3D vertical NAND flash (a.k.a. V-NAND). V-NAND stacks layers of flash cells on top of one another in order to increase flash capacity instead of shrinking feature sizes. SSDs with 32-layer V-NAND are already available on the market.

Then, the speaker discussed two topics that his team was working on. The first topic concerned multi-streaming SSDs, which allow upper-level software to specify multiple (currently up to 8) I/O streams. For each stream, a multi-streaming SSD will monitor the stream's I/O pattern and optimize Flash management tasks accordingly. Therefore, upper level software can use the streams to segregate data that have different properties. Experimental study with popular NoSQL engines have shown significant gains in terms of steady-state SSD performance and device lifetime. Multi-streaming has been accepted by the SSD industry standard. Products with the multi-streaming feature will soon appear on the

market. The second part of the talk focused on Intelligent SSDs, which support In-Storage Computing (ISC) and run user-written codes within the SSDs. ISC exploits the internal bandwidth of flash chips and the processing cycles within the SSD. Operations (e.g., filtering) that process a large amount of data but produce only a small result set can significantly benefit from ISC. The speaker presented the promising results of several case studies using database and data analytics workloads.

3. RESEARCH PAPERS

The second session featured four paper presentations:

- (1) “Scalable and Efficient Spatial Data Management on Multi-Core CPU and GPU Clusters: A Preliminary Implementation based on Impala” by Simin You, Jianting Zhang, and Le Gruenwald. This paper integrates a spatial indexing and query processing engine with Cloudera Impala. The engine was previously developed by the authors. It supports both multi-core processors and GPUs. The integration effort modifies the Impala frontend to support spatial query syntax and implements a spatial data management module in impala that invokes the spatial query engine as a shared library. Compared to a traditional single-core technique, the resulting solution has achieved orders of magnitude of speedups on a high-end GPU-equipped workstation. In addition, the resulting solution has shown high efficiency and good scalability on a 10-node Amazon EC2 cluster equipped with multi-core CPUs and GPUs.
- (2) “Optimizing CPU Cache Performance for Pregel-Like Graph Computation” by Songjie Niu and Shimin Chen. In-memory graph computation systems have been used to support many important applications, such as PageRank on the web graph and social network analysis. This paper studies the CPU cache performance of graph computation. To facilitate the study, the authors implemented a graph computation system, called GraphLite, in C/C++ based on the description of Pregel. The paper analyzes the CPU cache behavior of the internal data structures and operations of graph computation, then exploits CPU cache prefetching techniques to improve the cache performance. Preliminary experiments on a real machine show that the proposed technique can achieve about 2x speedups for PageRank computation.
- (3) “Query Processing on Low-Energy Many-Core Processors” by Annett Ungethüm, Dirk Habich, Tomas Karnagel, Wolfgang Lehner, Nils Asmussen, Marcus Völp, Benedikt Nöthen and Gerhard Fettweis. This paper studies the techniques to implement query processing on Tomahawk, a low-energy heterogeneous multiprocessor system-on-a-chip. Tomahawk

consists of an application core, a larger number of processing elements (PEs), and a core manager controller, all connected through a network-on-chip. The paper investigates two APIs of Tomahawk, a CUDA-like programming interface and a micro-kernel interface. Analysis and experimental evaluation show that the former has a simpler programming interface, while the latter allows programmers more controls to place different operators on different PEs, thereby avoiding unnecessary data transfers incurred by the former API.

- (4) “Flash-aware Index Scan in PostgreSQL” by Da-som Hwang, Woon-hak Kang, and Sang-won Lee. This paper combines sorted index scan with parallel I/Os on SSDs to improve index scan performance. Sorted index scan is a technique to sort the record IDs from the range scan of the secondary index before retrieving the records. Parallel I/Os can take advantage of the internal parallelism of SSDs. The paper modifies the index scan implementation in PostgreSQL. Preliminary experimental results show that the optimization achieves dramatic improvements for index scans, and the optimized index scan can be faster than a full table scan even with 100% selectivity.

4. CONCLUSION

Over 30 people attended the half-day meeting. Both the keynote talk and the research paper presentations stimulated interesting questions and discussions. This shows significant interests of the ICDE community in the theme of HardBD: exploiting new hardware technologies for efficient Big Data management. From the discussions, it is clear that there are many open research problems, and both academia and industry have been actively working in this area.

5. ACKNOWLEDGMENT

Many people contributed to the successful organization of HardBD 2015. We would like to thank all the authors for their contributions, the PC members for excellent reviewing work, the keynote speaker for delivering an insightful speech, and the ICDE 2015 Workshops co-chairs for their guidance and support.

This workshop is partially supported by the CAS Hundred Talents program, State Key Laboratory of Computer Architecture (ICT, CAS), NSFC Innovation Research Group No. 61221062, and the National 863 Hightech Program (No. 2013AA013204).

6. REFERENCES

- [1] J. Gehrke, W. Lehner, K. Shim, S. K. Cha, and G. M. Lohman, editors. *31st IEEE International Conference on Data Engineering, ICDE 2015, Seoul, South Korea, April 13-17, 2015*. IEEE, 2015.