Distributed and Parallel Databases (2020) 38:767–769
https://doi.org/10.1007/s10619-020-07313-y

Introduction to the special issue on Self-managing and Hardware-Optimized Database Systems 2019

Shimin Chen[1] · Panos K. Chrysanthis[2] · Khuzaima Daudjee[3] · Meichun Hsu[4] · Mohammad Sadoghi[5]

Published online: 27 September 2020

During the past 40 years, data management systems have grown in scale, complexity, and variety. There have been novel extensions to RDBMS as well as fast evolution of big data systems, such as Key-Value stores, Document stores, Graph stores, Spark, MapReduce/Hadoop, Graph Computation Systems, and Data Stream Processing System. At the same time, hardware technology in processors, memory, storage, and networking is undergoing rapid changes, while administration and tuning of these systems has become very expensive, imposing new challenges for data management systems.

Two specialized IEEE ICDE workshops, namely SMDB (International Workshop on Self-Managing Database Systems) and HardBD&Active (Joint International Workshop on Big Data Management on Emerging Hardware and Data Management on Virtualized Active Systems), have provided a forum to examine the above system-related challenges from different angles. The SMDB workshop focuses on providing autonomic or self-* features in database systems to support complex

✉ Shimin Chen
chensm@ict.ac.cn

Panos K. Chrysanthis
panos@cs.pitt.edu

Khuzaima Daudjee
kdaudjee@uwaterloo.ca

Meichun Hsu
meichun.hsu@oracle.com

Mohammad Sadoghi
msadoghi@ucdavis.edu

1 Institute of Computing Technology, Chinese Academy of Sciences, Beijing, China

2 University of Pittsburgh, Pittsburgh, USA

3 University of Waterloo, Waterloo, Canada

4 Oracle Corporation, Redwood City, USA

5 UC Davis, Davis, USA

administrative tasks, while the HardBD&Active workshop is interested in exploiting hardware technologies for efficient data management. We are pleased to present the first special issue of DAPD entitled "Self-Managing and Hardware-Optimized Database Systems" that features the best contributions in SMDB 2019 and HardBD&Active 2019 workshops.

In the following, we provide a brief overview of the five papers in this special issue.

In "Revisiting Hash Join on Graphics Processors: A Decade Later", Paul et al. revisit the major GPU hash join algorithms in the last decade with an emphasis on the various GPU features introduced over time and exploited in the algorithms. They present a comprehensive evaluation using different generations of GPUs to understand the impact of the GPU features on hash join performance. In addition, the authors propose techniques to cope with data skews and high match rates in GPU hash joins, and compare GPU hash joins with state-of-the-art CPU implementation in the evaluation.

In "Self-Driving Database Systems: A Conceptual Approach", Kossmann and Schlosser present a component-based framework for self-driving database systems that use reusable and exchangeable components to handle subproblems in the context of system self-management. They propose a linear programming based algorithm to derive an efficient tuning order automatically for optimizing multiple mutually dependent features, e.g., index selection and compression configurations. The applicability and scalability of the approach is demonstrated with case studies.

In "A Gray-Box Modeling Methodology for Runtime Prediction of Apache Spark Jobs", Al-Sayeh et al. studies the challenging problem of predicting Spark job runtimes, which depend on numerous factors such as the input data cardinalities, allocated resources, and configuration parameters. The authors propose a gray-box modeling methodology that comprises a white-box model for predicting RDD cardinalities, and a black-box model for predicting the runtime of each task for given RDD cardinalities. Experimental evaluation shows that this methodology can accurately predict runtimes of Spark jobs, such as word count, PageRank, and K-means.

In "AutoTuning Memory Management in Data Analytics", Kunjir studies an important problem of auto-tuning the memory allocation for applications running on modern distributed data processing systems. The author proposes a framework called Guided Bayesian Optimization (GBO) that uses a white-box model that separates good configurations from bad ones, as a guide during the black-box Bayesian Optimization exploration process used in many autonomous data system solutions. Experimental evaluation on Apache Spark using industry-standard benchmark applications shows that GBO can significantly speed up the optimization process.

In "Initial Experience with 3D XPoint Main Memory", Liu and Chen experiment with 3D XPoint main memory, the first commercially available main memory NVM solution targeting mainstream computer systems. They find that 3D XPoint conforms to many assumptions about NVM in previous simulation or emulation studies: it is 2–3X slower than DRAM, and persisting data from CPU cache to 3D XPoint incurs drastic overhead. Moreover, 3D XPoint has a number of distinctive features: 3D XPoint's performance degrades as the data set size increases, and the amount of modifications in a cache line does not impact its write latency.

Finally, we would like to thank all the authors who have contributed their papers for this issue. We own our sincere gratitude to the reviewers who contributed to assembling such a high-quality special issue. We are also indebted to the DAPD Journal Editors, editorial office, and the publishing and production teams for their assistance in preparation and publication of this issue.

Guest Editors:

Shimin Chen, Institute of Computing Technology, Chinese Academy of Sciences (lead editor)
Panos K. Chrysanthis, University of Pittsburgh
Khuzaima Daudjee, University of Waterloo
Meichun Hsu, Oracle Corporation
Mohammad Sadoghi, UC Davis