The VLDB Journal
https://doi.org/10.1007/s00778-024-00841-4

EDITORIAL

Special issue: modern hardware

Norman May[1] · **Spyros Blanas**[2] · **Danica Porobic**[3]

Accepted: 10 January 2024

The continued evolution of computing hardware and infrastructure imposes new challenges and bottlenecks to program performance. As a result, traditional database architectures that focus solely on I/O optimization increasingly fail to utilize hardware resources efficiently. Multi-core CPUs, GPUs, FPGAs, new memory and storage technologies (such as flash and nonvolatile memory), and low-power hardware impose a great challenge to optimizing database performance. Consequently, effectively utilizing the characteristics of modern hardware has become an important topic of database systems research.

In this special issue on modern hardware, we provide a snapshot of the research in this vibrant area. This issue follows a series of special issues on this topic including the special issue on modern hardware in 2016, and special sections with best papers of DaMoN, the International Workshop on Data Management on New Hardware.

This special issue followed an open call for submissions. Two papers in this issue are extended versions of papers from DaMoN 2022 and another contribution is an extended version of a paper from ICDE 2023. All the papers have additional new material compared to the workshop or conference version and underwent the VLDB Journal's normal reviewing process. The four papers that were finally selected for this special issue all went through a major revision between July and September 2023, then a minor revision between October and November 2023, before being accepted in December 2023. We next present a brief summary of the accepted papers.

✉ Norman May
norman.may@sap.com

Spyros Blanas
blanas.2@osu.edu

Danica Porobic
danica.porobic@oracle.com

[1] SAP SE, Walldorf, Germany

[2] Ohio State University, Columbus, OH, USA

[3] Oracle Switzerland, Zürich, Switzerland

The first paper presents a hybrid CPU-FPGA platform where work is distributed between CPU and FPGA based on heuristics. To enumerate triangles, the heuristic assigns lightweight intersections to the CPU, while computation-intensive intersections are assigned to the FPGA using a hash-based method. The resulting algorithm also supports triangle enumeration on dynamic graphs, i.e., graph updates are supported. The experiments reveal that this hybrid processing strategy outperforms a pure CPU-based alternative, and it is more energy efficient than selected GPU-based alternatives.

The second submission extends a prior DaMoN paper. It presents a buffer caching strategy that improves OLAP performance on high-bandwidth storage. Unlike conventional frequency-based replacement strategies, the paper considers the expected speedup for the caching decision. For example, heavy computation may hide the data access latency so that caching may not be needed. These ideas are integrated into a prototype of a query execution engine, and this platform was used for extensive experiments demonstrating the benefits of the proposed new caching strategy.

The third paper extends a DaMoN paper exploring almost latch-free data structures that provide scalability on today's multi-core processor architectures. Previously proposed latch-free approaches required complex designs specialized for a particular component of the database system. With the newly proposed epoch-protection version scheme (EPVS), authors propose a more practical approach where most of the accesses are latch-free with rare complex operations requiring latching. The proposed EPVS approach achieves performance comparable to the specialized latch-free approaches while being general enough to easily integrate inside a full featured data management system.

Finally, the fourth paper extends an ICDE paper studying indexing approaches tuned for a novel trend of memory disaggregation where CPUs access not only locally attached, but also remote memory through high speed interconnect. It proposed a LSM-based approach, dLSM, that matches various stages of computation, including compaction and

Published online: 22 February 2024

checkpointing, to available resources across different servers adaptively. In this way, it can achieve good speedups over the state-of-the-art indexing approaches, thus making it practical to use disaggregated memory effectively.

We would like to thank the program committee members of the recent DaMoN workshops and the VLDB Journal reviewers for their efforts in reviewing the papers that have made this special issue possible. We would also like to thank the authors for their efforts in producing significant new material and results for this special issue that address opportunities for improvement that were identified during the VLDB Journal review process.

As hardware architectures and devices evolve, the research at the intersection of modern hardware and data processing will continue to offer rich opportunities for innovations. We are grateful to have been given the opportunity to produce a special issue on this important topic.