2nd Workshop on Novel Data Management Ideas on Heterogeneous (Co-)Processors (NoDMC)

Dirk Habich[1] David Broneske[2]

The objective of this one-day workshop is to explore the challenges and opportunities of data processing on existing and future heterogeneous hardware architectures. On the one hand, today's processors are no longer mainly bound by the density and frequency of transistors, but by their power and heat budgets. The so-called "power wall" forces hardware suppliers to rely more on the design of specialized devices optimized for certain types of calculations, which results in an increasingly heterogeneous processor landscape. On the other hand, memory and storage has seen an unprecedented change as well: novel and already commercially available techniques have blurred the traditional mental picture of a memory/storage hierarchy. For example, Non-Volatile RAM (NVRAM) is a prominent example to question the long-standing memory hierarchy reflected in almost all system-level applications. Moreover, very large caches, High-Bandwidth-Memory (HBM), Non-Uniform Memory Access (NUMA), or even remote-memory designs as well as extremely fast SSDs add to the heterogeneous portfolio of available memory/storage techniques. Therefore, to meet the performance requirements of the modern information society, tomorrow's database systems will have to exploit and embrace this increased heterogeneity of processor and memory technologies.

The purpose of this workshop is to assist with training and fostering a community of researchers and industry practitioners working on data processing issues on heterogeneous hardware systems. To this end, we want to provide a forum to discuss challenges, progress and directions, and to offer an environment for networking persons researching on related topics and fostering future collaborations. Especially in the view of the SPP 2037 on *Scalable Data Management for Future Hardware* and the SPP 2377 on *Disruptive Memory Technologies*, we want to strengthen collaborations beyond individual SPP projects by connecting them with other researchers. This workshop is co-organized by the GI-Arbeitskreis *Data Management on Modern Hardware*.

The scope of the workshop includes, but is not limited to:

[1] TU Dresden, dirk.habich@tu-dresden.de
[2] German Centre for Higher Education Research and Science Studies, broneske@dzhw.eu

doi:10.18420/BTW2023-40

- Applications of modern hardware in
  - data mining
  - data-intensive machine learning
  - query processing
  - sensor or stream processing
  - non-traditional applications (e.g., graph processing)
- Algorithms and data structures for efficient data processing on and across different (co-)processors or memory technologies
- Exploitation of specialized ASICs or specialized memories technologies (e.g., processing in memory (PIM))
- Efficient memory management, data placement and data transfer strategies in heterogeneous systems
- Energy efficiency in heterogeneous hardware environments
- Programming models and hardware abstraction mechanisms for writing data-intensive algorithms on heterogeneous hardware
- Query optimization, cost estimation and operator placement strategies for heterogeneous hardware
- Transaction processing in heterogeneous systems

With the given scope of the workshop, we are happy to announce a great program. The workshop starts with a keynote by David F. Bacon working at Google Research, who is the leading architect of the Spanner storage engine. From the submissions, we were able to accept five technical papers as well as four extended abstracts. The corresponding talks are organized in two sessions according to the topics of *Advances in Storage, Memory, and Network Technologies* in Session 1 and *Advances in (Co-)Processing Technologies* in Session 2. The first talk in the first session is by Baumstark et al. who investigate the capabilities of processing-in-memory technologies for table scans. Afterward, El-Shaikh et al. present how to store information using DNA-based storage systems. The third technical paper by Lutsch et al. focuses on the performance of SGX for machine learning workloads. Next, the extended abstract of Benson et al. gives lessons learned of using persistent memory under CXL. The last talk in this session by Geyer et al. discusses benefits and drawbacks of CXL for heterogeneous cloud architectures.

In Session 2, four papers in the area of (co-)processor acceleration are presented. The first talk by Damme and Boehm is an extended abstract to a CIDR paper presenting an architecture for exploiting heterogeneous processors for data science applications. Afterward, Hahn et al. present an FPGA parser and an according parser generator for the Apache Avro

format – a semi-structured data format used in stream processing applications. The third talk is by Schuhknecht and Islam, who benchmark heterogeneous multi-core CPUs running multiple queries at a time in parallel. The workshop program closes with the presentation by Fett et al. who investigate the performance of matrix multiplication under different compressed formats and overflow handling on GPUs.

Last but not least, we like to thank everyone who contributed to this workshop, in particular, the authors, the reviewers, the BTW team, and all participants.

PC Chairs

- David Broneske (DZHW)
- Dirk Habich (TU Dresden)

Steering Committee

- Wolfgang Lehner (TU Dresden)
- Gunter Saake (University of Magdeburg)
- Kai-Uwe Sattler (TU Ilmenau)

Program Committee

- Sebastian Breß (Snowflake)
- David Broneske (DZHW)
- Patrick Damme (TU Berlin)
- Philipp Götze (SAP SE)
- Dirk Habich (TU Dresden)
- Tilmann Rabl (HPI Potsdam)
- Hannes Rauhe (SAP SE)
- Horst Schirmeier (TU Dresden)
- Knut Stolze (IBM Germany)
- Annett Ungethüm (Universität Hamburg)
- Stefan Wildermann (Friedrich-Alexander Universität Erlangen-Nürnberg)
- Steffen Zeuch (DFKI und TU Berlin)