

Special Issue on the Best Papers of DaMoN 2019

Ashraf Abounaga

the date of receipt and acceptance should be inserted later

The annual workshop on Data Management on New Hardware (DaMoN) has established itself as a key publication venue in the database community. The topics covered by the workshop range from benchmarking database systems on new hardware to hardware-conscious database techniques to new hardware designs optimized for database systems.

The fifteenth DaMoN workshop was held on July 1, 2019, in Amsterdam, in conjunction with the SIGMOD Conference. In recent years, DaMoN has become one of the best-attended pre-conference workshops at SIGMOD, and 2019 was no exception. The workshop was co-chaired by Thomas Neumann and Ken Salem, and the program contained thirteen research papers.

This special “Best of DaMoN” section contains extended versions of three papers from the workshop. The papers were selected by the workshop co-chairs and invited to submit to this special section. All the papers have at least 30% new material compared to the workshop version and underwent the VLDB Journal’s normal reviewing process, with some of the original DaMoN reviewers and new reviewers who were not part of the DaMoN program committee.

The first paper, titled “Building Blocks for Persistent Memory”, presents a benchmarking study of Intel’s Optane DC Persistent Memory Modules. This is significant since persistent memory (PMem) promises to bridge the gap between DRAM and SSDs, and this paper presents the first comprehensive evaluation of PMem on real hardware. Based on the evaluation results, the paper proposes four building blocks for database systems or other applications using PMem.

The second paper, titled “VIP: A SIMD-Vectorized Analytical Query Engine”, presents an analytical query engine called VIP (for Vectorization, Interpretation, and Partitioning). VIP uses pre-compiled column-oriented data-parallel sub-operators with SIMD-implementations. VIP represents an interesting point in the design space of analytical query engines since it does not use code generation or operator pipelining. The paper shows that these design choices result in better performance in memory-bound systems.

The third paper, titled “Faster & Strong: String Dictionary Compression using Sampling and Fast Vectorized Decompression”, improves upon a state-of-the-art dictionary-based string compression technique. The paper proposes a fast compression technique with a reduced memory footprint. In addition, it uses SIMD-based optimizations to accelerate access to the compressed data. This is important since string compression is widely used in database systems, and the string dictionaries can represent a large part of the memory consumption of a database system (up to 40% in one real-world system presented in the paper).

My thanks to the DaMoN co-chairs and program committee and to the VLDB Journal reviewers. Their efforts in selecting and reviewing these papers made this special section possible.

Ashraf Abounaga
Guest Editor of the Special Section