



Special issue on best papers of DaMoN 2018

Kenneth Salem¹

Published online: 23 December 2019
© Springer-Verlag GmbH Germany, part of Springer Nature 2019

The annual workshop on Data Management on New Hardware, known as DaMoN, has become a leading, high-profile venue for research on database systems. As the workshop's title suggests, its particular focus is the boundary at which database systems meet the evolving landscape of computing hardware. How can database systems exploit new hardware features, and how can hardware evolve to better meet the needs of data management?

The 2018 workshop was the fourteenth in the DaMoN series. It was held on 11 June in Houston, Texas, in conjunction with ACM's annual SIGMOD International Conference on Management of Data. Hundreds of attendees heard presentations of ten papers as well as a keynote talk about persistent memory from Intel's Andy Rudoff. The 2018 program also introduced "Fresh Thinking" talks: mini-keynotes from new researchers, including Spyros Blanas (Ohio State), Jana Giceva (Imperial College), and Pinar Tözün (IT University of Copenhagen). The workshop was co-chaired by Ken Salem (University of Waterloo) and Wolfgang Lehner (TU Dresden), and papers were reviewed by a program committee consisting of fourteen leading international academic and industrial researchers.

This special "Best of DaMoN" section is a result of new agreement between the DaMoN steering committee and the editors-in-chief of the VLDB Journal. To select the papers for the special section, I drew up an initial shortlist based on the reviews and program committee discussions of the DaMoN workshop papers. This shortlist was then whittled down to three papers through consultations with three senior members of the DaMoN program committee, each of whom considered all of the shortlisted papers. The authors of these three papers were invited to submit extended versions of their DaMoN

work for this special section. The extended submissions went through the VLDB Journal's normal peer-reviewing process, including new reviewers who were not part of the DaMoN program committee.

The first paper, from TU Munich and CWI in Amsterdam, looks at the use of SIMD instructions in database query pipelines. SIMD instructions allow many tuples to be processed in parallel by each instruction, but parallelism can suffer because of control flow divergence: Different tuples may require different handling. This paper shows how to combat this problem by refilling SIMD lanes that have been idled because of divergence.

The second paper, from TU Dresden, considers the problem of persistent memory failures, which can lead to data loss and downtime. This problem can be combated by replicating data, and this paper proposes efficient mechanisms for doing so.

The third paper, from TU Ilmenau, considers the problem of how to exploit high-bandwidth memory on many core processors for database join processing. It provides advice on how best to do so, backed by extensive experiments.

I'm pleased to present the three papers in this special section as representatives of the excellent work that appeared at DaMoN 2018. Thanks to the DaMoN program committee and to the additional VLDB Journal reviewers for their work in selecting and reviewing these papers.

Kenneth Salem
Guest Editor of the Special Section
Program Co-Chair of DaMoN 2018

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

✉ Kenneth Salem
ken.salem@uwaterloo.ca

¹ University of Waterloo, Waterloo, Canada