

Special issue on best papers of VLDB 2015

Chen Li¹ · Volker Markl²

Published online: 15 December 2016
© Springer-Verlag Berlin Heidelberg 2016

The VLDB 2015 conference, which was held in Kohala Coast, Hawaii, received 710 high-quality submissions, with 139 papers accepted for presentation at the conference. Including rollover papers from PVLDB 2014, this resulted in a program with 160 research papers in total. Out of this research program, we invited seven outstanding contributions to provide an extended version of their work to the VLDB Journal in a special issue on best papers of VLDB 2015. After the review process, six submissions were accepted.

The paper “Dissociation and Propagation for Approximate Lifted Inference with Standard Relational Database Management Systems” by Wolfgang Gatterbauer and Dan Suciu focuses on the approximate evaluation of conjunctive queries. The authors propose a method for probabilistic inference, where a query is evaluated entirely in the database engine by evaluating a fixed number of query plans, each providing an upper bound on the true probability, then taking their minimum. The authors show that by considering schema information it is possible to generate a minimal set of query plans within a database engine, thus computing the propagation score efficiently. For self-join-free conjunctive queries, there is only a single plan, in which case the authors show that the query evaluation is in PTIME.

The paper “Fast and Scalable Inequality Joins” by Zuhair Khayyat, William Lucia, Meghna Singh, Mourad Ouzzani, Paolo Papotti, Jorge-Arnulfo Quiane-Ruiz, Nan Tang, and Panos Kalnis describes inequality join algorithms based on

sorted arrays and space-efficient bit arrays. The authors also introduce a simple method to estimate the selectivity of inequality joins in order to optimize conjunctive queries and multi-way joins. The authors provide implementations and evaluations both in a centralized and in a distributed setting.

The paper “Resource Bricolage and Resource Selection for Parallel Database Systems” by Jiexing Li, Jeffrey F. Naughton, and Rimma V. Nehme was a runner up for the VLDB 2015 best paper. The authors investigate the challenges of running parallel database systems in an environment with heterogeneous resources, in particular the problem of unequal load across machines resulting in underutilization of powerful machines as well as the problem of selecting the right set of heterogeneous machines in order to process a given workload. The authors formalize the two problems. They employ mixed-integer programs to search for the optimal resource selection and verify their approach with an extensive experimental study.

The paper “Multi-objective parametric query optimization” by Immanuel Trummer and Christoph Koch presents parametric query optimization techniques to handle multiple objectives. The authors formalize the multiple objective parametric query optimization problem and present two algorithms to address it. They prove that both algorithms cover the space of relevant query plans and provide an experimental performance analysis.

The paper “Incremental Knowledge Base Construction Using DeepDive” by Christopher De Sa, Alex Ratner, Christopher Ré, Jaeho Shin, Feiran Wang, Sen Wu, and Ce Zhang gives a comprehensive description of DeepDive, a system that combines database and machine learning ideas to help develop inference-based knowledge-base construction systems. The authors develop techniques to support the iterative process of knowledge-base construction to incrementally produce inference results. They also study the trade-off

✉ Chen Li
chenli@gmail.com
Volker Markl
volker.markl@tu-berlin.de

¹ University of California, Irvine, CA, USA

² Technische Universität Berlin (TU Berlin), Berlin, Germany

space of the proposed techniques and provide an overall evaluation.

The paper “Order Indexes: Supporting Highly Dynamic Hierarchical Data in Relational Main-Memory Database Systems” by Jan Finis, Robert Brunel, Alfons Kemper, Thomas Neumann, Norman May, and Franz Faerber investigates the problem of how to maintain and query hierarchical data in a relational database system in the presence of complex, possibly skewed structural updates at high data rates. The authors propose the concept of order indexes, a dynamic representation of the nested intervals encoding hierarchical data. They also give an evaluation, which shows the update efficiency

and robustness of the approach for highly dynamic workloads.

We believe that these six papers, which were among the best publications coming out of VLDB 2015, represent a diverse and exciting set of innovative research contributions, which advance the state of the art in data management systems and beyond.

Chen Li and Volker Markl
VLDB 2015 PC Chairs and PVLDB Volume 8 Editors in Chief