

Editor's Notes

Welcome to the June 2017 issue of the ACM SIGMOD Record!

First of all, we welcome the new officers of the SIGMOD Executive Committee, including Juliana Freire as the Chair, Ihab Francis Ilyas as the Vice-Chair, and Fatma Ozcan as the Treasurer. The new SIGMOD Executive Committee takes office in June 2017.

The first column of this issue is the Database Principles column, featuring an article by Barceló, Píeris, and Romero on semantic optimization of Conjunctive Queries (CQs). As one of the most fundamental classes of database queries, CQs correspond to select-project-join queries in relational algebra. The database theory community has identified the classes of CQs that can be evaluated in polynomial time, i.e., the CQs that admit a suitable tree decomposition of small width. Furthermore, semantic information about the data, in the form of constraints, can be used to enrich query optimization by guiding the query transformation process. This article considers such semantic query optimization for the tractable classes of CQs and answers questions including whether a set of constraints can be used to reformulate a CQ as one of small width, and if so, what is the cost of computing and evaluating such a reformulation. As such, the article provides us with a theoretical framework for reasoning about the efficiency of such query reformulations.

The Survey column features an article by Singh and Bawa on MapReduce-based spatial query processing approaches. Support of high performance queries on spatial data has been an important topic in database research. This article presents classification and analysis of recent spatial query processing approaches, implemented in the MapReduce framework, into two categories. The first category includes hierarchical index approaches and the second category includes key-value storage based index approaches, with the main differences lying in the way that a spatial index is implemented on the partitioned dataset.

The Systems and Prototypes column features an article on the ARCHIMEDES system for efficient query processing over probabilistic knowledge bases. Due to the uncertainty of information extraction algorithms and the limitations of human knowledge, current knowledge bases are still incomplete and uncertain. The ARCHIMEDES system addresses query processing in probabilistic knowledge bases with three key technical components: (a) knowledge expansion derives implicit knowledge from knowledge bases using large rule sets; (b) query-driven inference improves inference performance by focusing MCMC on the query variables; (c) the system further leverages unified data- and graph-parallel computation to improve performance.

The Distinguished Profiles column features Beng Chin Ooi, Distinguished Professor of Computer Science at the National University of Singapore (NUS). Beng Chin is the recipient of the 2009 SIGMOD Contributions Award, and he is an IEEE and ACM Fellow and Fellow of Singapore National Academy of Science. In this interview, Beng Chin talks about his research vision (in 2011 when this interview took place), how he built up a successful research group and advised students at NUS, how he interacted with the Chinese database research community, his entrepreneur experience, and finally his work ethic and favorite pastimes.

The Reports column features a report on the third workshop on Algorithms and Systems for MapReduce and Beyond (BeyondMR'16) held in conjunction with the 2016 SIGMOD conference in San Francisco, California, USA. The goal of the workshop was to bring together researchers and practitioners to explore algorithms, computational models, architectures, languages and interfaces for systems that need large-scale parallelization and systems designed to support efficient parallelization and fault tolerance. The workshop program featured two very well attended invited talks by Ion Stoica from AMPLab, University of California Berkeley and Carlos Guestrin from the University of Washington. The program also included 5 regular and 5 short papers on specialized programming and data management systems based on MapReduce and extensions, graph processing systems, and data-intensive workflow and dataflow systems.

On behalf of the SIGMOD Record Editorial board, I hope that you enjoy reading the June 2017 issue of the SIGMOD Record!

Your submissions to the SIGMOD Record are welcome via the submission site:

<http://sigmod.hosting.acm.org/record>

Prior to submission, please read the Editorial Policy on the SIGMOD Record's website:

<https://sigmodrecord.org>

Yanlei Diao

June 2017

Past SIGMOD Record Editors:

Ioana Manolescu (2009-2013)	Alexandros Labrinidis (2007-2009)	Mario Nascimento (2005-2007)
Ling Liu (2000-2004)	Michael Franklin (1996-2000)	Jennifer Widom (1995-1996)
Arie Segev (1989-1995)	Margaret H. Dunham (1986-1988)	Jon D. Clark (1984-1985)
Thomas J. Cook (1981-1983)	Douglas S. Kerr (1976-1978)	Randall Rustin (1974-1975)
Daniel O'Connell (1971-1973)	Harrison R. Morse (1969)	