

Editor's Notes

Welcome to the September 2015 issue of the ACM SIGMOD Record!

This issue opens with a letter from the SIGMOD Executive Committee, summarizing some of the major changes implemented by the SIGMOD Executive Committee and other SIGMOD news of the last two years. The major changes range from the revisions of the SIGMOD conference, to the new special issues of the SIGMOD Record, to the development of a new “sigmod.org” website, and to the new rewards in the SIGMOD community.

This issue continues with a Database Principles article by Wenfei Fan, which provides an overview of recent advances in the study of data quality, from theory to practice. The fundamental research questions covered in this article include data consistency, data deduplication, information completeness, data currency, and data accuracy. Following the theoretical questions, the article surveys a range of practical techniques for profiling (discovery of data quality rules), cleaning (error detection and data repairing), and matching (data deduplication). The article closes by pointing out challenges introduced by big data to data quality management.

The Research and Vision Articles Column features a vision article, by Gawlick et al., on “Mastering Situation Awareness”. This article is motivated by the observation that in applications such as cloud operation and customer care, situation awareness requires a system to support the management of data, knowledge, processes, and other services such as social networking in an integrated way. To this end, the authors propose a model, called KIDS (Knowledge Intensive Data-processing System), which enables the development and maintenance of situation-aware applications in a declarative and therefore economical manner. The article also outlines a list of opportunities and challenges for the database community, including a theoretical foundation based on category theory, time dimension support, performance and scalability.

The Surveys Column features two articles. The first survey, by Yin et al., examines “Robust Query Optimization Methods With Respect to Estimation Errors.” Query optimization has been a key component of a database management system (DBMS) since the design of the System-R optimizer. Effective query optimization, however, depends on accurate estimation of data and query related statistics, which remains as a technical challenge. This survey particularly focuses on a recent notion of “robust query optimization”, which aims to find a plan (or several plans) whose query execution time differs from that of an optimal plan by only a small fraction. Toward this goal, the article surveys a wide range of techniques, including cardinality injection, dynamic re-optimization, deferred plan selection, tuple routing, etc., under a set of criteria for comparison. As such, the article provides a good introduction to the topic of robust query optimization and characterizes a large set of recent techniques developed for this purpose.

The second survey by Kim and Lee focuses on “Community Detection in Multi-Layer Graphs”. The goal of community detection (graph clustering) is to partition vertices in a complex graph into densely-connected components, called communities. In recent applications, however, an entity is often associated with multiple aspects of relationships—capturing such multiple aspects of interactions requires multi-layer graph clustering, where each graph represents an aspect of the interactions. This article presents a large number of multi-layer graph datasets used in recent studies, surveys and compares recent relevant techniques, and suggests future directions for research.

The Distinguished Profiles column features Rick Cattell. Rick spent over 20 years at Sun Microsystems, where he was involved in many technology innovations that we take for granted today, such as ODBC, JDBC, and J2EE. Rick was one of Sun's first Distinguished Engineers, and his dissertation on compiler technology won the ACM Dissertation Award. In this interview, Rick shared his experiences and thoughts regarding software patents, standards, and performance, scalability, and availability of database systems.

This issue includes three event reports. The first article, by Pedersen, reports on the International Workshop on Energy Data Management (EnDM 2014), co-located with EDBT/ICDE 2014. The workshop, targeting at novel schemes for large-scale energy data processing, featured 5 research papers on modeling, semantics, and analytics of energy data. The second article, by Chen et al., reports on the International Workshop on Big Data Management on Emerging Hardware (HardDB 2015), co-located with ICDE 2015. The workshop included a keynote speech and four research papers that exploit new hardware trends, such as multi-core CPU and GPU clusters, for big data analytics. The third article, by Castellanos et al., reports on the International Workshop on Business Intelligence for the Real-time Enterprise (BIRTE 2014), co-located with VLDB 2014. The workshop attracted a large audience with a full program including a keynote speech by C. Mohan from IBM Research, two industrial talks from Microsoft and HP Labs, respectively, and four research papers on various aspects of real-time business analytics.

Finally, this issue closes with the call for participation for ICDE 2016, to be held in Helsinki, Finland in May 2016.

On behalf of the SIGMOD Record Editorial board, I hope that you all enjoy reading the September 2015 issue of the SIGMOD Record!

Your submissions to the 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:

<http://www.sigmod.org/publications/sigmod-record/sigmod-record-editorial-policy>

Yanlei Diao

September 2015

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)	