Editor's Notes

Welcome to the December 2016 issue of the ACM SIGMOD Record!

This issue opens with the 2016 Dagstuhl Report that outlines some of the most important research directions in the area of Principles of Data Management (PDM). In April 2016, a group of researchers in the PDM area joined in a workshop at the Dagstuhl Castle in Germany, which was organized jointly by the Executive Committee of the ACM Symposium on Principles of Database Systems (PODS) and the Council of the International Conference on Database Theory (ICDT). As summarized in this report, the workshop identified research challenges for PDM around seven core themes, namely, *Managing Data at Scale*, *Multi-model Data*, *Uncertain Information*, *Knowledge-enriched Data*, *Data Management and Machine Learning*, *Process and Data*, and *Ethics and Data Management*. In addition, the workshop highlighted two accelerating trends: the increasing embrace of neighboring disciplines, including especially Machine Learning, Statistics, Probability, and Verification; and the increased focus on obtaining positive results to enable the use of mathematically-based insights in practical settings. These trends are expected to continue in the years to come.

The Database Principles column continues with an article by Koutris and Suciu on a formal analysis of multiway join processing in massively parallel systems. The article introduces a theoretical model, called the MPC (Massively Parallel Computation) model, which characterizes the number of synchronization points and the maximum load per machine as observed in popular Hadoop and Spark systems. Using the MPC model the article shows the design of novel algorithms for multiway join and theoretical results that prove their optimality through tight lower bounds.

The Vision column features two articles. The first article, by Kamat and Nandi, studies variance implementations in real-world database systems given their increased importance in big data analytics. The article reports that some major database systems use a representation that suffers from floating point precision loss, then reviews literature on variance calculation in both the statistics and database communities, and finally gives recommendations on implementing variance functions in various query processing settings. The second article, by Vartak et al., studies visualization recommendation systems. With the advent of large, high-dimensional datasets and significant interest in data science, there is a growing need for tools that can support rapid visual analysis. This article presents the vision of a new class of visualization systems that can automatically identify and interactively recommend visualizations relevant to an analytical task. The article describes the key requirements for such a visualization recommendation system as well as the challenges in realizing this vision, and finally presents several approaches to address the challenges.

The Distinguished Profiles column features Rick Hull, a distinguished researcher, and Stratos Idreos, a recent recipient of the SIGMOD Jim Gray Dissertation Award. Rick Hull is currently a researcher at IBM. Before joining IBM, he was a professor at the University of Southern California for many years and managed a research group at Bell Labs. Rick is an ACM Fellow and a coauthor of the classic database theory book, "Foundations of Databases." In this interview, Rick talks about his work experiences in different institutes and highlights a few research topics such as artifact-centric business process models and cognitive computing. The second interview features Stratos Idreos, the 2011 recipient of the SIGMOD Jim Gray Dissertation Award. Stratos' thesis addresses Database Cracking, that is, auto-tuning of database kernels. Stratos is currently an Assistant Professor at the Harvard University.

The Reports column features a report on the First International Workshop on Reproducible Open Science (RepScience2016), which was organized in conjunction with the 20th edition of the International Conference on Theory and Practice of Digital Libraries. The goal of the workshop was to provide a forum for constructively exploring foundational, organizational, and systemic challenges towards the implementation of Open Science publishing principles. The workshop brought together skills and experiences in the form of invited talks and paper presentations, focusing on the definition and establishment of the next generation scientific communication ecosystem.

On behalf of the SIGMOD Record Editorial board, I hope that you enjoy reading the December 2016 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

December 2016

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)