

Editor's Notes

Welcome to the June 2014 issue of the ACM SIGMOD Record!

The issue opens with a Database Principles article by Ameloot on recent theoretical work on declarative networking. This article addresses three important aspects of declarative networking: (a) *coordination* where nodes of a cloud are trying to obtain a global consensus, and in particular, the CALM conjecture by Hellerstein, (b) *correctness* of distributed computation, including decidability results, and (c) *languages* in declarative networking and *declarative semantics* for such languages, including the CRON conjecture by Hellerstein. The article concludes by outlining several directions for future work. This article presents a timely review of the recent theoretical work on declarative networking and offers a great introduction to this topic for the interested reader.

The Research and Vision Articles column features three articles. First, Morton, Balazinska, Grossman, Kosara, and Mackinlay, present an analysis of usage patterns of Many Eyes and Tableau Public, two popular Web-based, collaborative visual analytics systems. The analysis explores a number of primary dimensions of online visual analytics including the types of users, how users collaborate and interact, and how they analyze single datasets versus multiple data sources. The results of this study offer valuable information and insights towards building online visual analytics systems in the future. The second article, by Torres, Galante, and Pimenta, addresses the issue that today's Object-Relational Mappings are often platform dependent and deeply embedded in the code, which is difficult to read, understand, or evolve. This article presents ENORM, a notation that extends class models representing all the essential mappings and does so in a platform-independent manner. The third article, by Lin, Chang, and Chao, addresses keyword search queries over XML documents which may contain arbitrary combinations of AND, OR, and NOT operators. This article presents the concept of *valid Smallest Least Common Ancestors (SLCAs)* as query results, which eliminates erroneous results returned by previous algorithms, as well as an efficient algorithm to process such queries.

The Systems and Prototypes Column features the Medusa system developed by Zhong and He, which is a parallel graph processing system on graphics processors (GPUs). The core design of Medusa is to enable developers to leverage the massive parallelism and other hardware features of GPUs by writing sequential C/C++ code for a small set of APIs. The runtime system of Medusa then automatically executes the user-defined APIs in parallel on the GPU, with a number of optimizations based on the architecture features of GPUs and the characteristics of graph applications. A case study in social network analysis shows how Medusa improves both the coding productivity and the performance of graph operations.

In the Research Centers column, Haas, Cefkin, Kieliszewski, Plouffe, Roth, et al. describe the design and activities of the IBM Research Accelerated Discovery Lab. The lab is built on an analytics cloud environment with a unique software system that supports the process of discovery, facilitating collaboration and fostering insight. While many other groups have or are creating institutes that focus in one way or another in data-driven discovery, the IBM Research Accelerated Discovery Lab is unique in its emphasis on supporting the overall discovery process and the focus on understanding, from a social science perspective, how discovery happens and how it may be accelerated. These aspects are illustrated through a diversity of analytic and systems research projects that span disciplines and institutions, as well as through a study of the practice of discovery, with the goal to use the findings to better enable and accelerate discovery.

This issue features three event reports. Koutrika, Lakshmanan, Riedewald, and Stefanidis report on the First International Workshop on Exploratory Search in Databases and the Web (ExploreDB 2014), held in conjunction with EDBT/ICDT 2014. The report highlights the keynote by Prof. Keim on "Exploring Big

Data using Visual Analytics”, outlines a collection of six papers on various topics ranging from data exploration with structured database queries to search and ranking on the Web, and summarizes the panel discussion at the workshop. The report concludes by pointing out a number of research directions in areas including databases, Web search, multimedia exploration. In the second article, Ma, Meng, and Wang report on the Sixth International Workshop on Cloud Data Management (CloudDB 2014), co-located with ICDE 2014. The report covers the two keynotes speeches, “Building Big Data Processing Systems under a New Computing Model” by Prof. Xiaodong Zhang and “Multi-faceted Classification of Big Data Uses and Proposed Architecture Integrating High Performance Computing and the Apache Stack” by Prof. Fox, as well as 10 research papers covering a wide range of topics from Quality of Service, to Query Processing, System Architecture, and Benchmarks. The report closes by pointing out some open problems including big data management in the cloud and cloud data security and privacy. The third article, by Manghi, Bolikowski, Houssos, and Schirrwagen, reports on the First Workshop on Linking and Contextualizing Publications and Datasets, held in conjunction with the 3rd International Conference on Theory and Practice of Digital Libraries (TPDL). This workshop addresses the need to interlink and contextualize datasets and scientific publications in sectors of scholarly communication and digital libraries. The report covers two keynote speeches and 10 research papers on topics including dataset contextualization, interlinking datasets and publications, and representing and visualizing datasets. The report concludes by outlining main considerations and future issues with respect to publications and datasets.

On behalf of the SIGMOD Record Editorial board, I hope that you will all enjoy reading the June 2014 issue of the SIGMOD Record.

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

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

Prior to submitting, you are encouraged to read the Editorial Policy on the SIGMOD Record’s Web site (<http://www.sigmod.org/publications/sigmod-record/sigmod-record-editorial-policy>).

Yanlei Diao

June 2014

Past SIGMOD Record Editors:

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