

Editor's Notes

From Athens at the end of our annual SIGMOD conference, welcome to the March 2011 issue of the ACM SIGMOD Record!

The first paper of the issue by Ling, Chang and Chao focuses on the efficiency of answering keyword search queries in XML databases. Keyword search is a tempting alternative to structural queries, which are perceived as relatively complex to write and require some knowledge of the XML document organization. Natural candidate answers to an XML keyword search queries are least common ancestors (LCAs) of all the query keywords, however, some of them are intuitively more information-rich and more specific than the others. The algorithms provided in this paper allows speeding up the task of identifying the strongest contributors, that is, the most interesting nodes to return as results of an XML keyword search.

The Systems and Prototypes article by Chen, Kannan, Madhavan and Halevy describes SchemR, a visual search engine and repository for schemas. The motivation for building schema repositories is to simplify the task of designing and deploying a data management application, by enabling would-be user to reuse schemas that others have designed for the same or similar applications. SchemR users can search the repository of schemas by using keywords and/or a partial schema; a comprehensive GUI allows users to navigate the returned results, zoom through the results etc. Importantly, SchemR is part of the OpenII framework, and thus available for other researchers to use.

This issue's Distinguished Database Profile features Dennis Shasha commenting on his many interests, curiosities, stories and insights. As anyone who has read his books (and in particular "Out of their minds") knows, beyond being a great scientist, Dennis is a fascinating storyteller, as this column fully shows. I have had many opportunities to enjoy discussing with Dennis. In INRIA in 2000-2001, when we were office mate, Dennis was presenting us the (then) very new idea of privacy-preserving joins; on a highway crossing primary forests in Brasil, we talked about biology experiment design and how various parts of the world evolved over the centuries. In this column Dennis shares his stories on database tuning, queries by humming, puzzles, and much more.

The Open Forum article by Aumüller and Rahm describes an analysis of recent-years publications in SIGMOD, VLDB, VLDBJ and TODS. Based on DBLP records, the authors analyze the evolution of affiliations of authors by institutions, countries and continents, the structure of international collaborations etc. Following up on a similar SIGMOD Record by M. Nascimento in 2003, this paper provides a quite detailed image of the last 10 years' evolutions in publications in these important venues.

Finally, the issue features three event reports.

The Workshop on Data-Intensive Software Management and Mining was co-located with the ACM CIKM 2010 conference. The main idea behind the gathering lies in four main aspects of the large set of projects whose sources are available today: they can be seen as a web-scale repository, continuously evolving, containing complex-structure items, and forming the basis of a social network of interactions.

The Dagstuhl workshop on Bi-Directional Transformations is a follow-up of a previous conference on the same topic, held in 2008. The workshop reunited researchers from Programming Languages, Graph Transformations, Software Engineering and Database communities in discussions around models for transforming complex (typically graph-structured) data, and their correctness, completeness, possible symmetry, and supported consistency checking.

The Dagstuhl Advanced School on Data Exchange, Data Integration and Streams (DEIS 2010) has been a remarkable innovation over the usual format of a summer school, allowing for much more intensive participation of the attendant PhD students and young researchers. Each of the 22 participants was assigned a topic by the organizers, had to prepare a 45-minute presentation of the topic and had the opportunity to interact with one organizer while preparing and improving the presentation. This interactive mode was highly appreciated by organizers and participants alike.

I take advantage of these notes to signal you the almost complete renewal of the SIGMOD advisory board, as well as the end of Lisa Singh's tenure as the SIGMOD conference coordinator. Lisa's outstanding service ended, Sihem Amer-Yahia takes over the duty of eliciting, handling and shepherding bids for organizing our yearly conference. The recipients of the 2011 SIGMOD awards have also been added to the list of colleagues we honor in each issue.

Your contributions to the Record are welcome via the RECESS submission site (<http://db.cs.pitt.edu/recess>). Prior to submitting, be sure to peruse the Editorial Policy on the SIGMOD Record's Web site (<http://www.sigmod.org/publications/sigmod-record/sigmod-record-editorial-policy>).

Ioana Manolescu

June 2011

Past SIGMOD Record Editors:

Harrison R. Morse (1969)
Daniel O'Connell (1971 – 1973)
Randall Rustin (1975)
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)