

Editor's Notes

Welcome to the March 2012 issue of the ACM SIGMOD Record! This is a rich issue with many interesting contributions.

This issue opens with an article by Chang on the optimization of XML twig pattern queries featuring full-text search predicates. The optimization of XML twig pattern queries has received significant attention in the past; at the same time, full-text search is very useful in text-rich XML databases, and a significant extension of XQuery specifically deals with full-text search. The article explores a space of optimization techniques exploiting together the structural and text constraints of the query, proposing a cost model and optimization heuristics.

In the Database Principles column, the paper by Peter Wood surveys query languages for graph databases, a topic on which attention is renewed due to the popularity of important applications such as Semantic Web graphs of RDF data, and social network graph analysis. Query languages are analyzed from the viewpoint of features (expressive power) and then under the angle of their associated algorithmic complexity.

Two surveys appear in this issue. The first one, by Dustdar, Pichler, Savenkov and Truong, attempts to systematize an area at the confluence of data management and Web services, namely data services. The authors argue for a new approach in modeling, describing and integrating distributed and heterogeneous data sources, by means of quality-aware data services, and discuss issues they raise, such as metadata management, data service publication, and the selection of the “best” services to be used for a given integration task. The survey by Mami and Bellahsene considers the problem of recommending views to materialize in order to improve the performance of a given query workload; the authors analyze and classify many algorithms from the existing literature from the angle of the general view selection framework they use (data cube lattice, ANR-OR view graph etc.) as well as from the cost constraint perspective.

The issue includes two Distinguished Profiles in Databases. The first one features David Lomet, whose outstanding career has been distinguished among others by two ACM SIGMOD Best Paper awards, as well as being an ACM Fellow and an IEEE Fellow. The column is remarkable in many respects, of which I'll just select two. First, beyond the many successful projects, successes, we also learn of decisions and technical views that David came to regret later on, which I find reassuring for the rest of us mortals! Second, the interview smoothly takes you quite advanced points of detail concerning transactions, isolation, latches, all the way up to flash disks and cloud computing. Our second distinguished columnist is Catriel Beeri, also an ACM Fellow. His interview blends the story of his research career between Jerusalem, Toronto, and Princeton, advising Moshe Vardi and sharing a room with Phil Bernstein, with beautiful stories of ideas, such as data dependencies and the chase. Read it to find out the full implications and connections between the data dependency theory, data integration, data-exchange, view-based query evaluation and more – and also the possibly lesser-known music lover side of Catriel!

The Research Centers column features a presentation of the Institute for the Management of Information Systems, within the “Athena” Research center, headed by Timos Sellis. The article is dense with descriptions of the IMIS' research areas: Data lifecycles, Geoinformatics, Web of Data, Biological Data Management, Privacy and more!

The Industry Perspective column features an article on the new SQL:2011 standard, issued in December 2011. Although these days a lot of attention is given to systems claiming to be not (or not only!) SQL, one should still not ignore the elephant in the living room – namely, the still-dominant industry standard

language. The article by Zemke outlines the extensions brought by the new standard, such as: enhancement of collection types, pipelined data manipulations and more.

Finally, in the Open Forum column, Grossniklaus and Maier outline the curriculum in cloud-based data management at Portland State University, in Oregon, USA. The authors detail their choices of courses, and student projects, and recommended readings. In this very fast-changing landscape of technology, this presentation is timely as many University departments are pondering the organization of similar courses.

This year's SIGMOD conference will be held early, on May 22-24 in Scottsdale, Arizona, as usual jointly with PODS and several interesting workshops. I'll be there and looking forward to any feedback you may have on the Record's overall organization and content! We will also be looking for new editors, contact me at the conference if you are interested.

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

March 2012

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)