

Editor's Notes

Welcome to the September 2011 issue of the ACM SIGMOD Record. Thanks to our energetic editors, and in particular to the newcomers (Pablo Barceló for Database Principles, Vanessa Braganholo for Distinguished Profiles in Database Management, and Glenn Pauley for Industrial Perspectives), all columns are represented!

I take the opportunity to thank Leonid Libkin for his remarkable service as a Database Principles editor over twelve years. Leonid has stepped down and the column is in the good hands of Pablo now. Thank you, Leonid!

The issue starts with a Database Principles column by Cali, Gottlob, Lukasiewicz and Pieris. This article is very timely and most welcome, coming at a short interval after the remarkable related invited lecture by Georg Gottlob at ICDE 2011. Organized as a short survey, it informs our database-oriented readership on recent advancements in querying data with attached knowledge, or ontology-enhanced databases. The article outlines variants of Datalog +/- languages, and provides their complexity results.

The paper by Buneman, Cheney, Lindler and Mueller describes their Database Wiki platform for collaborative data management. The tool seeks to combine the best of both the relational database and Wiki worlds: structured, robust data management as provided by databases, and collaborative, flexible, unstructured as supported by Wikis. The platform is particularly tailored to support and simplify collaborative curation of data, by providing annotations, versioning, and provenance tracking.

This issue's Distinguished Profiles in Databases column features two interviews. Meral Özsoyoğlu is the first to discuss with Marianne her research work, life lessons and more. In her interview, Meral first discusses genealogical data management, which many readers will think they are reasonably acquainted with; in contrast, most of us know much less about metabolic pathways databases or metabolic networks, which she does a very good job at explaining here. Working across scientific disciplines is never easy, and Meral discusses the particular case of bioinformatics, pleading that the challenges and potential success deserve the risk-taking.

The second interview of the column features Divesh Srivastava. Beyond his amazingly productive career, Divesh is a great source of stories and insights. As anyone who has ever talked to him at a conference knows, he is also a seemingly endless source of new problems: after patiently listening to my VLDB demo in a freezing room at Cairo's Mena House Oberoi hotel, he playfully offered four to six other problems to work on! It was reassuring (of sorts) to learn years later that he would do the same to some of my PhD students. Read the interview to find out the role of AT&T as a communication company, what real life databases look like, how the large IIT-originating database community has formed, and more.

The report on the Data Management Group at NEC Research Labs highlights the activity of a recent team, created in 2009, and whose research work mainly focused on cloud-based data management. Under the umbrella of their CloudDB project, the authors survey their research on elastic OLTP workloads, based on the concept of transaction class which allows trading consistency for elasticity, one of the most attractive features of cloud computing. Further, the report outlines work on resource and workload management in heterogeneous clouds, and on cloud-supported mobile applications.

In the Industrial Perspectives column, Inkster, Zukowski, and Boncz present an overview of the integration of the VectorWise column store into the Ingres relational database system. In particular, the authors describe the approaches taken to the store and the server catalog, the integration of query processing, and the techniques used to convey execution plans between the systems and the return of results back to the client.

Two articles appear in this issue's Open Forum column. The one by Ameloot, Marx, Martens, Neven and Van Vees nicely summarizes the main facts, tendencies, authors and topics present in the 30 years of existence of the ACM Principles of Databases (PODS) conference. Do not miss out the lyrics of the famous song performed at the conference anniversary in Athens! Also, the authors had compiled a great set of graphs summarizing the PODS conference activity. Since all of them did not fare well in print, the paper points you to the associated Web site.

The second paper by Badia and Lemire addresses the authors' perceived insufficiency of database design methodology and techniques, at least as we have learned them and possibly teach them in school today! The authors offer a compelling set of anecdotes showing how these techniques are unused, misused and insufficient, and seek to open a conversation on these topics within the community. We found the paper controversial, but interesting food for thought.

Last but not least, the SIGMOD 2012 call for papers closes the issue: the place is Scottsdale, Arizona, and the research paper deadline is November 1st! Of course, all information can be found on the conference Web site (<http://www.sigmod.org/2012/index.shtml>).

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

October 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)