

Editor's Notes

Welcome to the December 2008 issue of SIGMOD Record. We begin the issue with a welcome message from Yanniss Ioannidis to SIGMOD members in general and to student members in particular.

In lieu of regular articles, this issue has a **Special Section on Managing Information Extraction**, edited by AnHai Doan, Luis Gravano, Raghu Ramakrishnan, and Shivakumar Vaithyanathan. The field, fueled by the explosion of unstructured data on the Web, has become increasingly important in recent years, receiving attention from multiple communities (AI, DB, WWW, KDD, Semantic Web, and IR). The nine papers included in this special section highlight the wide range of Information Extraction problems; they are properly presented in the introduction to the special issue. Many thanks to AnHai, Luis, Raghu, and Shivakumar for their hard work in putting together this special section, and especially to AnHai who took care of all the details with the manuscript submission system.

The **Database Principles Column** (edited by Leonid Libkin) features two articles. The first article, by Kimelfeld and Sagiv, surveys recent results on modeling and querying probabilistic XML data. The second article, by Koch, proposes a query algebra for probabilistic databases.

The **Distinguished Profiles in Data Management Column** (edited by Marianne Winslett) features an interview of Surajit Chaudhuri, who is a research area manager at Microsoft Research, an ACM Fellow, and has received the SIGMOD Contributions Award in 2004. Read Surajit's interview to find out (among many other things) how data mining led him to self-tuning databases and about life as a research manager.

We continue with an article in the **Research Centers Column** (edited by Ugur Cetintemel) about the ETH Zurich Systems Group and Enterprise Computing Center and, in particular, their research projects across three technology trends: multicore, virtualization, and cloud computing.

Next is the **Open Forum Column**, which is meant to provide a forum for members of the broader data management community to present (meta-)ideas about non-technical issues and challenges of interest to the entire community. In this issue, we have a fantastic paper by Graham Cormode, about *how NOT to review a paper*, i.e., how to avoid being an adversarial reviewer. I really enjoyed reading this article, I am sure you will too.

We continue with five articles in the **Event Reports Column** (edited by Brian Cooper). First is the *Report on the First European Conference on Software Architecture (ECSA 2007)*, written by Cuesta and Marcos. Second is the *Report on International Workshop on Privacy and Anonymity in the Information Society (PAIS 2008)*, written by Xiong, Truta, and Fotouhi. Third is the *Report on the IFIP WG5.8 International Workshop on Enterprise Interoperability (IWEI 2008)*, written by van Sinderen, Johnson, and Kutvonen. Fourth is the *Report on the First Workshop on Very Large Digital Libraries (VLDDL 2008)*, written by Manghi, Pagano, and Zezula. Finally, fifth is the *Report on the First Workshop on Transforming and Weaving Ontologies in Model Driven Engineering (TWOMDE 2008)*, written by Silva, Pan, Assmann, and Herinksson.

We close the issue with multiple **Calls for Papers/Participation**:

- SIGMOD 2009 Workshops: DaMoN, DBTest, IDAR, MobiDE, and WebDB,
- ER 2009: 28th International Conference on Conceptual Modeling, and
- WSDM 2009: 2nd ACM International Conference on Web Search and Data Mining

Alexandros Labrinidis
January 2009