

Editor's Notes

Welcome to the March 2008 issue of SIGMOD Record. We begin the issue by a regular **article** on *distributed databases and peer-to-peer databases* (by Bonifati, Chrysanthis, Ouksel, and Sattler). The article examines the emergence of P2P databases and compares them to distributed, federated, and multi-databases; the authors characterize the differences among these types of databases by distinguishing between P2P-centric and db-centric features.

We continue with an article in the **Surveys Column** (edited by Cesar Galindo-Legaria), on *Querying Encrypted XML Documents* (by Unay and Gundem). The survey compares the various algorithms for query processing over encrypted XML documents, which is becoming increasingly important for enabling the “database as a service” paradigm.

Next we have an article on the **Systems and Prototypes Column** (edited by Magdalena Balazinska), about the *BP-Mon*, which is a system to support Query-Based Monitoring of BPEL Business Processes. The article written by Beerli, Eyal, Milo, and Pilberg was invited after the authors' VLDB 2007 demo.

The **Distinguished Profiles in Data Management Column** (edited by Marianne Winslett) features an interview of Serge Abiteboul who is a senior researcher at INRIA and the manager of the Gemo Database Group. Read Serge Abiteboul's interview to find out (among many other things) about building a research group in Europe, why systems papers should not have to include measurements, etc.

We continue with an article in the **Research Centers Column** (edited by Ugur Cetintemel), about *Data Management Projects at Google* (by Halevy, Madhavan, Muthukrishnan, Cafarella, Chang, Fikes, Hsieh, and Lerner). The article highlights the group's research with special emphasis on the Deep Web and on large-scale data management projects.

We continue with the inaugural article of the **Open Forum Column**. The column 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. The article in this issue written by Manolescu, Afanasiev, Arion, Dittrich, Manegold, Polyzotis, Schnaitter, Senellart, Zoupanos, and Shasha, describes the experiences and feedback collected from the experimental repeatability process of the SIGMOD 2008 conference. This is a prime example of an article for the Open Forum Column (that even includes graphs!). I will be (mildly) editing the articles submitted in this column; if you have ideas about an article, please contact me to discuss it further.

Next we have six (yes, six!) articles in the **Event Reports Column** (edited by Brian Cooper).

- First is the *Report on the Third International Workshop on Computer Vision meets Databases (CVDB 2007)* which was held in June 2007, together with SIGMOD 2007.
- Second is the *Report on the Databases and Web 2.0 Panel at VLDB 2007* (by Amer-Yahia, Halevy, Alonso, Kossmann, Markl, Doan, and Weikum).
- Third is the *Report on the First International Workshop on Mining Graphs and Complex Structures (MGCS 2007)* which was held in October 2007, together with ICDM 2007.
- Fourth is the *Report on the Sixth ACM Workshop on Privacy in the Electronic Society (WPES 2007)*, which was held in October 2007, together with the ACM CCS Conference.
- Fifth is the *Report on the Tenth ACM International Workshop on Data Warehousing and OLAP (DOLAP 2007)*, which was held together with CIKM 2007.

- Sixth is the *Report on the Principles of Provenance Workshop*, which was held in November 2007.

We close the issue with two very important **Calls for Participation**: the *2008 SIGMOD/PODS Conference*, to be held in Vancouver, Canada, July 9-12, 2008 and the *Tribute to Honor Jim Gray*, which will be held on May 31, 2008 at UC Berkeley.

Alexandros Labrinidis
April 2008