

Editor's Notes

Welcome to the March 2010 issue of the ACM SIGMOD Record. It is a pleasure and an honor to write my first Editor's Notes for Record.

We begin the issue with a regular article by Hellerstein. The article is a companion to his keynote talk at the ACM PODS conference this year. The article is based on the previous works carried at Berkeley on distributed and parallel processing, and in particular applied to the networking domain. It addresses the challenges and opportunities of declarative database languages, and in particular Datalog, for parallelizing computations. While the content has first been presented in the theory-oriented setting of PODS, the article balances theoretical aspects with many practical applications and implementation issues. The importance taken by parallel and distributed processing in today's CS industry makes this article very timely.

The second regular article is by Barbieri, Braga, Ceri, Della Valle and Grossniklaus. It presents C-SPARQL, an extension to the SPARQL query language for querying streams of RDF data, as well as applications and future research opportunities. C-SPARQL allows querying RDF data both from static repositories and from streams, enabling continuous reasoning in contexts such as sensor networks or social semantic data.

The last regular article of the issue proposes declarative data-driven coordination (D3C) and is authored by Kot, Gupta, Roy, Gehrke and Koch. The work is motivated by the need to be able to specify, in a high-level fashion, complex yet flexible coordination mechanisms, such as "I want to book a seat in the same flight as my friend" and have such specifications compiled in lower-level transactions to be executed by a DBMS.

Two articles appear in the **Surveys Column** (edited by Cesar Galindo-Legaria). First, Khare, An and Song analyze and compare for the first time techniques used to automatically understand Web interfaces. Extracting data from the "Deep Web" requires the capacity to automatically fill in forms in order to get content, and understanding the forms is crucial in order to provide appropriate input. The authors set a classification framework to analyze works in this area, and illustrate it by analyzing 10 existing approaches. Second, Drosou and Pitoura consider diversification of answers in search and recommender systems. Diversification here is considered from the angles of content, novelty and coverage. The authors review definitions of diversity and concrete algorithms for achieving it.

The **Systems and Prototypes Column** (edited by Magdalena Balazinska) features one article by Liu, Mihaylov, Bao, Jacob, Ives, Loo and Guha. The article describes the SmartCIS environment for integrating physical and virtual environments, and follows the demo given by the authors at SIGMOD 2009. The SmartCIS prototype is based on the Aspen project, which extends the data integration formalisms to a setting of distributed streams.

The issue also features two workshop reports. The first is authored by Benedikt, Florescu, Gardner, Guerrini, Mesiti and Waller and summarizes the workshop on *Updates for XML* which was held in conjunction with the EDBT 2010 conference in Lausanne. The topics addressed in the workshop considered updates on XML and other data formats, such as RDF and probabilistic data. The second reports on the *Cloud Data Management Workshop* held next to the ACM CIKM 2009 conference in Hong Kong. The report witnesses of the interest raised by cloud computing within the data management community and concludes that (at the time of the workshop) much remained to be done in order to realize that potential.

Our issue closes with the announcements of DMSN, the Seventh International Workshop on Data Management for Sensor Networks, and XLDB, the 4th Extremely Large Database Conference, to be held in Menlo Park in October 2010.

This issue appears with some delay that we hope to absorb until the end of the year, by publishing the next three issues at short intervals in order to get closer to the normal publication rhythm. The SIGMOD Executive Committee continues consulting over the new editorial policy of the SIGMOD Record; we expect to publicly announce it in the fall.

Ioana Manolescu
August 2010

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)