Developments at ACM *TODS*

Richard Snodgrass

rts@cs.arizona.edu

The March 2005 issue of *TODS* has eight papers invited from the SIGMOD and PODS'2003 conferences. These papers are significantly extended versions of the conference papers, allowing the authors to refine and elaborate without the strictures of a twelve-page limit.

The first four papers in that issue were invited from the SIGMOD conference; the last four papers were invited from the PODS conference. Each went through the normal rigorous review process.

Alon Halevy was Program Chair for SIGMOD'2003; Tova Milo was Program Chair for PODS'2003. It is interesting that the first author of the first paper in this special issue is Tova. I should emphasize that while Tova helped select the PODS papers to invite, she had no involvement in selecting the SIGMOD papers to invite. It just turned out that one of the four SIGMOD papers selected was co-authored by her.

The June 2005 issue is almost complete; it will have at least seven papers (available now on the Upcoming Issues page¹). *TODS* continues to grow: the first two issues of 2005 contain more papers than the first two issues of 2004, and more than the number of papers that appeared in all of 2002. The full story is on the *TODS* web site².

I have appointed six new Associate Editors, bringing the complete Editorial Board to nineteen.

Jan Chomicki's research interests are in logical foundations of databases. Specific topics include: database integrity, data integration, data models, and query languages. His current projects involve query answering in inconsistent databases and preference queries.

Heikki Mannila works in data mining, algorithms and bioinformatics.

Raghu Ramakrishnan's current research is in two broad areas. In the EDAM project, he is working on data mining problems with driving applications in environmental monitoring, physics simulations, and e-commerce. In the CICADA project, he is working together with researchers from Microsoft Research on extending SQL to allow applications to specify when (potentially out of date) copies of data can be used.

Arnie Rosenthal's research is in the areas of data security and data sharing. In particular, he tries to align the technologies with feasible practices for data-owning organizations.

Sunita Sarawagi's research interests span several fields including databases, data mining, machine learning and statistics. Currently she is investigating the deployment of learning-based techniques for solving various data integration and cleaning tasks.

Dan Suciu works on applying formal theory to novel and difficult data management tasks. His past work was on various aspects of managing semistructured data, including query languages, compression, query processing and type inference, while his recent work focused on data security and on querying unreliable and inconsistent data sources.

All six are internationally-known scholars in the field of database systems and are well known also to the database community through their past service. In addition, they hail from three different continents.

I'm gratified that they are willing to help *TODS* continue to improve.

¹<http://www.acm.org/tods/Upcoming.html>

²<http://www.acm.org/tods/TurnaroundTime.html>