Foreword

This fourth issue of TODS Vol. 37, December 2012, has nine articles, five of which are invited.

A total of nine invited articles from four ACM database conferences, namely, the International Conference on Management of Data/Principles of Database Systems (SIGMOD/PODS 2011, Athens, Greece), and the International Conference on Database Theory/Extending Database Technology (ICDT/EDBT 2011, Uppsala, Sweden) were selected for publication in TODS. Four of these articles were published earlier in the third issue of TODS Vol. 37, and this issue contains the remaining five.

The first three invited articles in this issue are selected from PODS 2011. The article “Maximizing Conjunctive Views in Deletion Propagation” by Benny Kimelfeld, Jan Vondrak, and Ryan Willams studies the problem of propagating deletions of tuples in conjunctive views to database relations in such a way as to maximize the size of the updated view. The authors present nontrivial and elegant solutions. The second article, “Determining the Currency of Data” by Wenfei Fan, Floris Geerts, and Jef Wijsen, addresses the problem of reasoning about data currency in a data model in which data values participate in a partial order, stating that an item is more current than another. The authors propose a novel framework which leads to a natural modeling of real currency constraints and analyzes the complexity of several problems on querying the currency of data. The third article from PODS 2011, “Worst-Case I/O-Efficient Skyline Algorithms” by Cheng Sheng and Yufei Tao, studies the problem of the k -dimensional skyline in the external memory model and presents the first worst-case optimal algorithm for the skyline problem.

The next two articles in this issue are selected from SIGMOD 2011. The article “SCALLA: A Platform for Scalable One-Pass Analytics using MapReduce” by Boduo Li, Edward Mazur, Yanlei Diao, Andrew McGregor, and Prashant Shenoy presents an analysis of Hadoop and extends MapReduce to enable higher performance through hash-based group-by operations which also enable incremental computation for one-pass analytics. The second article, “WHAM: A High-Throughput Sequence Alignment Method” by Yinan Li, Jignesh Patel, and Allison Terrel, studies the problem of matching genome sequences efficiently and proposes a method designed by exploiting a hash index built on the reference sequence.

The first of the remaining four articles in this issue, “Finding Alternative Shortest Paths in Spatial Networks” by Kexin Xie, Ke Deng, Shuo Shang, Xiofang Zhou, and Kai Zheng, proposes a space-efficient algorithm to find alternative shortest paths in spatial network databases utilizing shortest-path quadrees. In the article “On Provenance Minimization”, Yael Amsterdamer, Daniel Deutch, Tova Milo, and Val Tannen study the part of provenance information (core) that depends on the query semantics and is invariant among equivalent queries and the problem of computing an equivalent query that realizes the provenance core for a given query. The next article, “Expressive Languages for Path Queries over Graph-Structured Data” by Pablo Barceló, Leonid Libkin, Anthony W. Lin, and Peter T. Wood, studies an extended class of conjunctive regular path queries, which can express relationships among paths, and provides detailed analysis of the complexity of such queries. Finally, the article “Partial Evaluation for Distributed XPath Query Processing and Beyond” by Gao Cong, Wenfei Fan, Anastasios Kementsietsidis, Jiangzhong Li, and Xianmin Liu studies the evaluation of XPath queries in a distributed setting based on partial evaluation.

Four additional invited articles from conferences in 2011 and selected for publication in TODS are already published in the third issue of TODS Vol. 37¹. These are “Revisiting Answering Tree Pattern Queries Using Views” by Junhu Wang and Jeffrey Yu, invited from EDBT 2011; “Efficient Reasoning about Data Trees via Integer Linear Programming” by Claire David, Leonid Libkin, and Tony Tan, invited from ICDT 2011; “Entangled Queries: Enabling Declarative Data-Driven Coordination” by Nitin Gupta, Lucja Kot, Sudip Roy, Gabriel Bender, Johannes Gehrke, and Christoph Koch, invited from SIGMOD 2011; and “Artifact Systems with Data Dependencies and Arithmetic Constraints” by Elio Damaggio, Alin Deutsch, and Victor Vianu, invited from ICDT 2011.

I would like to thank all the authors for their efforts to expand their articles according to the standards required for ACM TODS submissions. I thank the Program Committee Chairs Renee Miller of ACM SIGMOD 2011, Thomas Schwentick of PODS 2011, Tova Milo of ICDT 2011, and Natassa Ailamaki of EDBT 2011; the program committee members of all four conferences; and the conference and journal reviewers for their help in selecting these articles. Last but not least, I thank TODS Associate Editors, Divyakant Agrawal, Michael Bohlen, Graham Cormode, Susan Davidson, Wenfei Fan, Chris Jermaine, George Kollios, Tova Milo, Kenneth Ross, S. Sudarshan, Divesh Srivastava, and Yufei Tao for their dedicated editorial work on the articles in this issue and the papers invited from 2011 ACM database conferences.

Enjoy!

—Z. Meral Özsoyoğlu
Editor-in-Chief
 ACM Transactions on Database Systems

¹The first two of these four invited articles, published in TODS 37(3), were incorrectly listed as regular submissions in the foreword for TODS 37(3).