GUEST EDITORIAL

Special issue on best papers of VLDB 2012

**Z. Meral Özsoyoğlu · Uğur Çetintemel ·
Nilesh Dalvi · Hank Korth · Anthony Tung**

This special issue of VLDB journal contains the best papers from VLDB 2012, the 38th International Conference on Very Large Databases, 27–31 August, Istanbul, Turkey, with 726 participants. The conference received a record number of submissions as compared to previous years of VLDB.

- The Research Track received 659 submissions, of which 134 papers (20 %) were accepted. Out of 134 accepted papers, 78 were accepted in the first round, and 56 were accepted after they were resubmitted with revisions and a second round of reviewing.
- The Experiments and Analysis Track received 23 submissions, of which 8 papers (34.7 %) were accepted.
- The Industrial Track received 62 submissions, of which 16 papers (25.8 %) were accepted.
- The Demonstrations Track received 101 submissions, of which 34 (33.6 %) were accepted.

Z. M. Özsoyoğlu (✉)
Case Western Reserve University, Cleveland, OH, USA
e-mail: meral@case.edu

U. Çetintemel
Brown University, Providence, RI, USA
e-mail: ugur@cs.brown.edu

N. Dalvi
Facebook, Inc., Menlo Park, CA, USA
e-mail: nileshdalvi@gmail.com

H. Korth
Lehigh University, Bethlehem, PA, USA
e-mail: hfkorth@gmail.com

A. Tung
National University of Singapore, Singapore, Singapore
e-mail: akhtung@gmail.com

In addition, VLDB 2012 also included three keynote talks, pre- and post-conference workshops, tutorials, and panels.

The journal versions of seven best papers that appeared in the conference are included in this special issue of the VLDB Journal. These invited papers are substantially improved, revised, and extended as compared to their original conference versions and were accepted for publication after several rounds of journal-style reviewing. These papers cover a wide range of current database research topics, namely real-time graph data management, graph querying and indexing, social tagging of web data, probabilistic data consistency, adaptive indexing for transactional workloads, incremental view maintenance under high-rate updates, and storage provisioning.

The first paper of this special issue "Dense Subgraph Maintenance under Streaming Edge Weight Updates for Real-time Story Identification" (by Albert Angel, Nick Koudas, Nikos Sarkas, Divesh Srivastava, Michael Svendsen, and Srikanta Tirthapura) is also selected for the best paper award for VLDB 2012 conference. This paper addresses the problem of real-time story identification using social media data, such as daily blog posts and status updates, posted by millions of people around the globe. The main challenge the authors address is the efficient maintenance of dense subgraphs corresponding to groups of tightly coupled entities, under streaming updates of edge weights. Based on novel theoretical results on the amount of change resulting from a single edge weight update, authors present an algorithm for dense subgraph maintenance and demonstrate its performance with a thorough experimental evaluation.

The paper "An Expressive Framework and Efficient Algorithms for the Analysis of Collaborative Tagging" (by Mahashweta Das, Saravanan Thirumuruganathan, Sihem Amer-Yahia, Gautam Das, and Cong Yu) addresses the problem of collaborative tagging and develops a dual tagging

framework to explore the tagging behavior. Using similarity and diversity as measures, the authors define important analysis problems in this framework, demonstrate that they are NP-complete in the general case, and present efficient solutions.

The third paper, "Efficient Processing of K-Hop Reachability Queries" (by James Cheng, Zechao Shang, Hong Cheng, Haixun Wang, Jeffrey Xu Yu), presents a new index for efficient and scalable processing of k-hop reachability queries in large directed graphs. Efficiency and scalability of both the index construction and query processing are demonstrated using experiments on a range of datasets.

The paper "DBToaster: Higher-order Delta Processing for Dynamic, Frequently Fresh Views" (by Christoph Koch, Yanif Ahmad, Oliver Kennedy, Milos Nikolic, Andres Noetzli, Daniel Lupei, Amir Shaikhha) focuses on incremental view maintenance and presents a recursive solution for incrementally maintaining the view of an input query by utilizing discrete forward references, called delta queries, and then materializing the delta queries as views.

The paper "Quantifying Eventual Consistency with PBS" (by Peter Bailis, Shivaram Venkataraman, Michael J. Franklin, Joseph M. Hellerstein, and Ion Stoica) focuses on the tradeoffs between operation latency and the data consistency. By using the notion of Probabilistically Bounded Staleness (PBS), the paper quantitatively demonstrates that eventually consistent systems frequently return consistent data while offering significant benefits for latency. This explains why practitioners often prefer such systems.

The paper "Transactional Support for Adaptive Indexing" (by Goetz Graefe, Felix Halim, Stratos Idreos, Harumi Kuno, Stefan Manegold, and Bernhard Seeger) investigates the problem of concurrency control and recovery in the context of adaptive indexing where indexes are optimized incrementally as a side effect of query processing. Using detailed experimental analysis, the authors demonstrate several desirable properties of adaptive indexing, such as exploiting parallelism and maintaining adaptive properties under concurrent updates.

The final paper "Towards Cost-Effective Storage Provisioning for DBMS's" (by Ning Zhang, Junichi Tatemura, Jignesh Patel, and Hakan Hacigumus) presents a heuristic solution to the problem of provisioning data center resources in the I/O subsystem of a single node server for specific customer workloads while minimizing the total operating cost and respecting the existing service level agreements. Authors implemented this heuristic solution and demonstrated that it significantly reduces the total cost in various settings.

We thank the authors of invited papers for revising and extending their conference papers for submission to this special issue of VLDB journal. We also thank all the reviewers for detailed reviews with thoughtful and constructive comments. We hope you will find the papers in this issue full of inspiring research ideas as well as enjoyable to read.