

Special issue: best papers of VLDB 2010

Paolo Atzeni · Elisa Bertino ·
Masaru Kitsuregawa · Kian-Lee Tan

Published online: 21 February 2012
© Springer-Verlag 2012

This special issue of the VLDB Journal is dedicated to the best papers from the 36th International Conference on Very Large Data Bases, which took place on 13–17 September 2010 in Singapore.

The conference received 686 submissions overall. For the research tracks:

- The Core Database Technology Track received 280 submissions; of these, 19 were rejected without review because of formatting violations and 48 were accepted. The acceptance rate computed with respect to the reviewed papers is 18.4%.
- The Infrastructure for Information Systems Track received 215 submissions; of these, 11 were rejected without review because of formatting violations and 33 were accepted. The acceptance rate computed with respect to the reviewed papers is 16.1%.
- The Experimental and Analysis Track received 15 submissions and accepted 4.

P. Atzeni
Dipartimento di Informatica e Automazione,
Università Roma Tre, Rome, Italy
e-mail: atzeni@dia.uniroma3.it

E. Bertino (✉)
CS Department, CERIAS, Cyber Center (Discovery Park),
Purdue University, West Lafayette, IN, USA
e-mail: bertino@cs.purdue.edu

M. Kitsuregawa
Institute of Industrial Science,
The University of Tokyo, Tokyo, Japan
e-mail: kitsure@tkl.iis.u-tokyo.ac.jp

K.-L. Tan
Department of Computer Science, School of Computing,
National University of Singapore, Singapore 117417, Singapore
e-mail: tankl@comp.nus.edu.sg

In addition to the traditional program committees, papers were also reviewed through a “journal track”. A total of 119 submissions were received in this track during this year, of which 30 were accepted. In addition, there were 6 acceptances from other submission periods.

In addition, the Industrial Applications and Experience Track received 51 submissions, of which 12 (23.5%) were accepted. The Demonstrations Track received 89 submissions, of which 32 (36%) were accepted.

This special issue consists of journal versions of the five best papers among the papers accepted by the program committees of the Core Database Technology Track, the Infrastructure for Information Systems Track, and the Experimental and Analysis Track. The papers in the special issue cover a large spectrum of topics mirroring the breadth of the database field and showcasing the ongoing research activities and trends in this field. All papers have been significantly revised, improved, and extended beyond their initial version that appears in the proceedings of VLDB 2010. As customary for VLDB Journal, the papers went through several rounds of reviewing. We are grateful to the reviewers of these extended versions for their valuable comments and to the authors for their timely revisions of the papers.

The paper “HaLoop: Efficient Iterative Data Processing on Large Clusters”, by Yingyi Bu, Bill Howe, Magdalena Balazinska, and Michael Ernst, presents HaLoop, a modified version of the Hadoop MapReduce framework designed to serve many applications including data mining, web ranking, graph analysis, and model fitting, which require support for iterative programs. HaLoop shows interesting performances and retains the fault-tolerance properties of MapReduce through automatic cache-recovery and task re-execution.

The paper “MapMerge: Correlating Independent Schema Mappings” by Bogdan Alexe, Mauricio Hernandez, Lucian Popa, Wang-Chiew Tan studies mappings in data integration

or exchange and considers the issue of correlating multiple, independently designed schema mappings of smaller scope into larger schema mappings. This allows for a modular construction of complex mappings from smaller ones and enables a “divide-and-merge” paradigm for mapping creation.

The paper “Towards Certain Fixes with Editing Rules and Master Data” by Wenfei Fan, Jianzhong Li, Shuai Ma, Nan Tang, and Wenyuan Yu investigates the problem of repairing data in data monitoring. The paper introduces the notion of editing rules that can tell the attributes of an erroneous tuple to be fixed as well as the values that they should take from the master data. The paper studies the basic issues for reasoning about editing rules. It also presents an interactive framework and algorithm to find certain fixes. The experimental study on two real-life data sets shows that the proposed algorithm can fix errors in just two or three rounds of interactions.

The paper “Scalability of write-ahead logging on multi-core and multsocket hardware” by Ryan Johnson, Ippokratis Pandis, Radu Stoica, Manos Athanassoulis, and Anastasia Ailamaki examines the challenges to database system

scalability in multi-core and multi-socket hardware. In particular, it focuses on performance barriers due to logging and proposes Aether to eliminate lock contention, minimize context switches, and improve log buffer design. Furthermore, to manage the high communication costs between sockets in multsocket systems, AetherSMP extends Aether with a distributed log buffer design.

The paper “An Adaptive Updating Protocol for Reducing Moving Object Database Workload” by Su Chen, Beng Chin Ooi, and Zhenjie Zhang focuses on database systems for moving objects. Unlike most of the previous work that focuses on query performance and indexing techniques, this paper addresses the problem of the update model which is very critical for performance and, however, had not been investigated much. The paper proposes a generic and adaptive update model for moving object databases. The proposed model greatly reduces the number of update messages between the moving objects and the database server, thus reducing the overall workload of the system.

We hope that you will enjoy reading the papers in the special issue and get novel research ideas from these papers.