

# **Preface**

## **Special Issue: Selected Papers from ADBIS 2013 Satellite Events**

The 17th East-European Conference on Advances in Databases and Information Systems (ADBIS 2013) took place in Genoa, Italy, from September 1 to 4, 2013. The ADBIS series of conferences aims at providing a forum for the dissemination of research accomplishments and at promoting interaction and collaboration between the database and information systems research communities from Central and East European countries and the rest of the world. The ADBIS conferences provide an international platform for the presentation of research on database theory, development of advanced DBMS technologies, and their advanced applications. As such, ADBIS has meanwhile created a tradition: its 2013 edition continued the ADBIS series held in St. Petersburg (1997), Poznań (1998), Maribor (1999), Prague (2000), Vilnius (2001), Bratislava (2002), Dresden (2003), Budapest (2004), Tallinn (2005), Thessaloniki (2006), Varna (2007), Pori (2008), Riga (2009), Novi Sad (2010), Vienna (2011), Poznań (2012). The program of the 2013 edition included keynotes, research papers, and five satellite events, consisting of a Big Data special session, three thematic workshops, and a Doctoral Consortium.

In the present special issue, the extended and revised version of four papers out of the twenty-five papers presented at ADBIS 2013 Satellite Events are included. The papers cover various topics in large-scale data- and knowledge-centered systems, including GPU-accelerated database systems and GPU-based compression for large time series databases, design of parallel data warehouses, and schema matching.

The first two papers were presented at GID 2013 – Second International Workshop on GPUs in Databases.

The first paper, by Sebastian Breß, Max Heimel, Norbert Siegmund, Ladjel Belatreche, and Gunter Saake, is entitled “GPU-Accelerated Database Systems: Survey and Open Challenges” and explores the design space of GPU-accelerated database management systems. Based on the proposed survey, key properties and typical challenges of GPU-aware database architectures are presented and open research problems are formulated. Existing GPU-accelerated database management systems are also surveyed and their architectural properties classified, with a special emphasis on optimization issues. A reference architecture is finally proposed, indicating how GPU acceleration can be integrated in existing DBMSs.

The second paper is entitled “Compression Planner for Time Series Database with GPU Support” and is co-authored by Piotr Przymus and Krzysztof Kaczmarek. The paper exploits GPU in designing a compression planner for time series databases. Motivated by the fact that the growing volumes of time series data call for the definition of efficient and innovative processing approaches, the paper presents a novel compression method which is ultra fast and achieves the best possible compression ratio by composing several lightweight algorithms dynamically tuned for incoming data. The reported experimental results show that the proposed approach is a valid solution for data intensive monitoring and analytic systems.

The third paper, presented at SoBI 2013 – First International Workshop on Social Business Intelligence: Integrating Social Content in Decision Making, is entitled “A Global Paradigm for Designing Parallel Relational Data Warehouses in Distributed Environments” and is co-authored by Soumia Benkrid, Ladjel Bellatreche, and Alfredo Cuzzocrea. The paper proposes a novel methodology for designing a Parallel Relational Data Warehouse, in which, differently from other existing proposals, all the main design phases (i.e., fragmentation, allocation, and replication) are performed simultaneously, in a global fashion. The reported experimental results assess the performance of the proposed methodology against a well-known data warehouse benchmark.

The fourth paper, by Alsayed Algergawy, Seham Moawed, Amany Sarhan, Ali Eldosouky, and Gunter Saake, is entitled “Improving Clustering-Based Schema Matching Using Latent Semantic Indexing” and was presented at OAIS 2013 – Second International Workshop on Ontologies Meet Advanced Information Systems. The paper focuses on the identification of semantically corresponding elements across heterogeneous and large datasets and proposes a clustering-based matching algorithm which relies on a Latent Semantic Indexing-based approach in order to guarantee a high quality yet efficient process. Such an approach allows the retrieval of the conceptual meaning between clusters and the identification of the hidden semantic relationships among clusters elements. Experimental results are also reported, showing that the proposed approach is quite promising.

Several people have contributed to making this special issue possible. We are grateful to Abdelkader Hameurlain, Josef Küng, and Roland Wagner, Editors-in-Chief of the TLDKS journal, for accepting our proposal for this special issue. Special thanks are due to Gabriela Wagner, for her valuable assistance during the preparation of this special issue. We would also like to express our appreciation to the authors who submitted their manuscripts to be considered for this special issue. Last but not least, we would like to gratefully acknowledge the contribution of all the reviewers, who worked within a very tight schedule and whose detailed and constructive feedback to the authors contributed to improving the quality of the submitted manuscripts.

November 2014

Barbara Catania  
Giovanna Guerrini  
Themis Palpanas  
Jaroslav Pokorný  
Athena Vakali

Transactions on Large-Scale Data- and  
Knowledge-Centered Systems XV

Selected Papers from ADBIS 2013 Satellite Events

Hameurlain, A.; Küng, J.; Wagner, R.; Catania, B.;

Guerrini, G.; Palpanas, T.; Pokorný, J.; Vakali, A. (Eds.)

2014, IX, 125 p. 49 illus., Softcover

ISBN: 978-3-662-45760-3