
Editorial

Pedro Furtado* and Pedro Martins

Departamento de Engenharia Informática,
Faculdade de Ciências e Tecnologia,
Universidade de Coimbra,
Pólo II – Pinhal de Marrocos,
3030-290 Coimbra, Portugal
Email: pnf@dei.uc.pt
Email: pmom@dei.uc.pt
*Corresponding author

1 Introduction

The IEEE BigData Coimbra Satellite Session was a one-day satellite event of the 3rd International Congress on Big Data that happened from June 27 to July 2, 2014 in Anchorage, Alaska, USA. BigData@Coimbra was organised in Coimbra, Portugal, and it got together researchers and industry from the country and from other places in Europe to discuss together approaches for dealing efficiently with data of big volume, velocity and variety. We need to define architectures and approaches for dealing efficiently with BigData, cloud data and to learn more about applications and use cases. The IEEE BigData 2014 Coimbra Satellite Session focused on big data, cloud data management, and big data applications. The event featured research paper presentations, invited keynote speeches and use case presentations by industry. A work-in-progress (WiP) session was also organised, with discussions about ongoing promising work and poster presentations.

2 Contents of this special issue

This issue of the *IJBPM* contains a selection of research papers presented at the event.

The first paper from Antonio Rinaldi describes a novel framework to manage multimedia ontologies using a knowledge engineering approach, where the multimedia features are automatically extracted and multimedia ontologies are used to annotate and categorise images.

The second paper, by Jose Rodrigues et al., presents a new tool to identify and map speleothems by surveying cave chambers interiors. The chambers are surveyed using a terrestrial laser scanner to obtain point cloud datasets. The point cloud is used to generate four 3D-meshes to represent the surface model of the cave chamber with increasing levels of detail. These models are important for geomorphological studies or virtual visits to the cave. The paper we also presents an octree data structure implemented on the database that enables fast spatial queries.

The third paper, by Pedro Martins et al. describes an approach for automatically scaling the ETL process for freshness-preserving in high-rate data warehousing. The authors propose a universal data warehouse parallelisation solution, that is, an approach that enables the automatic scalability and freshness of any data warehouse and ETL.

The fourth paper, by Veronika Abramova et al., compares performance and scalability of SQL versus NoSQL systems under a varied set of workloads. CRUD operations are evaluated using the Yahoo! Cloud Serving Benchmark, while decision support operations are evaluated against MySQL Cluster and Hadoop with Hive.

The fifth paper, by Rafael Almeida et al., concentrates on the comparison of performance and scalability when executing the star schema benchmark on a cluster running Hadoop-Hive versus MySQL cluster, experimenting with demanding data sizes. The authors draw conclusions from the experiments concerning performance and scalability of each alternative and comparison between them.

The next three papers concern benchmarking setups and results dealing with diverse big data and scalability issues: clustering over SSB, a paper by Filipe Assunção et al.; big data scalable security, a paper by Joao Sousa et al.; Cassandra: what it does and what it does not and benchmarking, a paper by Melyssa Barata et al.

This special issue ends with an excellent survey of the field, by Alfredo Cuzzocrea. His work discusses the state-of-the-art on Data Warehousing and OLAP over Bigdata. Besides surveying the field, it also discusses open problems and future challenges, setting the stage for further research in this exciting subject.

As final remarks, we would like to thank every author for their contributions and participation in the event, and we sincerely hope readers will find the papers presented in this special issue very interesting and inspiring for further work on the subject.