Lecture Notes in Computer Science 8290

Commenced Publication in 1973
Founding and Former Series Editors:
Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Abdelkader Hameurlain Josef Küng
Roland Wagner Bernd Amann
Philippe Lamarre (Eds.)

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

Special Issue on Advanced Data Stream Management and Continuous Query Processing

Springer

ISSN 0302-9743 (LNCS) e-ISSN 1611-3349 (LNCS)
ISSN 1869-1994 (TLDKS)
ISBN 978-3-642-45268-0 e-ISBN 978-3-642-45269-7
DOI 10.1007/978-3-642-45269-7
Springer Heidelberg New York Dordrecht London

Library of Congress Control Number: Applied for

CR Subject Classification (1998): H.2.8, H.2, H.3, F.1, C.2, J.1

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India

Printed on acid-free paper

Springer is part of Springer Science+Business Media (www.springer.com)

Preface

An increasing number of applications continuously produce and consume high volumes of streaming data. Most of these applications are still developed from scratch and are rapidly overwhelmed by complex data management issues generated by the fast growth in the number of users, data items, and processing tasks. Data Stream Management Systems follow traditional database principles to facilitate the development of this kind of application via high-level programming interfaces and efficient data stream processing technologies built on declarative data stream models and languages, advanced query optimization algorithms, and new stream processing architectures.

This TLDKS Special Issue presents a representative selection of articles covering a wide range of important topics in advanced data stream management and processing of continuous queries. The book starts with application and development oriented contributions and ends with more theoretical results concerning data stream mining and analysis.

The first article "ASSIST: Access Controlled Ship Identification Stream" by Jianneng Cao et al. illustrates the application of algebraic rewriting techniques to avoid unauthorized access to streaming data. This article is particularly interesting because it also includes a real data stream management scenario involving maritime traffic control data of the port of Singapore shared among ship's captains, port authorities, shipping and insurance companies, cargo owners, and other stakeholders.

The design and development of applications mixing traditional databases and streaming data is a complex task for which it is crucial to have suitable tools. In their article "The HIT Model: Workflow-Aware Event Stream Monitoring" Olga Poppe et al. propose a workflow model combining continuous queries and timed automata to model complex real time reactive behavior of stream processing applications.

In a different direction, but still focussing on application development, the article "P-Bench: Benchmarking in Data-Centric Pervasive Application Development" by Sabina Surdu et al. characterizes, evaluates, and compares different event-centric and data-centric platforms for the development of applications in pervasive environments. Their work is a first step towards a benchmark for data stream management systems.

The amount of information published and disseminated in the form of data streams is constantly increasing and generates new challenges in the data analysis and knowledge discovery domain. In their article "Incremental Mining of Lag Patterns in Time Series Databases", Dhaval Patel et al. propose an adaptive and incremental approach to capture localized repeated associative relationships from multiple time series. In spite of the interest of such methods, for example for financial analysis, their use is currently limited by an excessive consumption

of resources. The incremental nature of their approach opens up the possibility a better scalability. Finally, Emmanuelle Anceaume et al. focussed their work entitled "On the Power of the Adversary to Solve the Node Sampling Problem" on the sampling of node identifier streams in unstructured peer-to-peer systems with malicious nodes. Node identifier streams constitute the keystone for many services (information dissemination, load balancing, data caching, etc.) and malicious nodes try to bias these streams, for example, to isolate some participant. The authors formally study the problem of obtaining uniform and fresh sampling strategies over input data streams generated by large open systems containing malicious participants.

The publication of this TLDKS Special Issue would not have been possible without the help of many people. First, we would like to thank all the authors who have submitted their articles and all the reviewers for their valuable efforts to ensure the high quality of this special edition. We also thank Abdelkader Hameurlain, Josef Küng, and Roland Wagner, Editors-in-Chief of the TLDKS journal, for their confidence and help. Finally, we are particularly grateful to Gabriela Wagner for her assistance in the preparation of these proceedings.

October 2013

Bernd Amann
Philippe Lamarre

Editorial Board

Reviewers

Table of Contents