

Commenced Publication in 1973

Founding and Former Series Editors:

Gerhard Goos, Juris Hartmanis, and Jan van Leeuwen

Editorial Board

David Hutchison

Lancaster University, UK

Takeo Kanade

Carnegie Mellon University, Pittsburgh, PA, USA

Josef Kittler

University of Surrey, Guildford, UK

Jon M. Kleinberg

Cornell University, Ithaca, NY, USA

Friedemann Mattern

ETH Zurich, Switzerland

John C. Mitchell

Stanford University, CA, USA

Moni Naor

Weizmann Institute of Science, Rehovot, Israel

Oscar Nierstrasz

University of Bern, Switzerland

C. Pandu Rangan

Indian Institute of Technology, Madras, India

Bernhard Steffen

University of Dortmund, Germany

Madhu Sudan

Massachusetts Institute of Technology, MA, USA

Demetri Terzopoulos

University of California, Los Angeles, CA, USA

Doug Tygar

University of California, Berkeley, CA, USA

Moshe Y. Vardi

Rice University, Houston, TX, USA

Gerhard Weikum

Max-Planck Institute of Computer Science, Saarbruecken, Germany

Christoph Bussler Malu Castellanos
Umesh Dayal Sham Navathe (Eds.)

Business Intelligence for the Real-Time Enterprises

First International Workshop, BIRTE 2006
Seoul, Korea, September 11, 2006
Revised Selected Papers

Volume Editors

Christoph Bussler
Cisco Systems Inc.
San Jose, CA 95134, USA
E-mail: chbussler@aol.com

Malu Castellanos
Hewlett-Packard
CA 94304, USA
E-mail: malu.castellanos@hp.com

Umesh Dayal
Hewlett-Packard
CA, 94304, USA
E-mail: umeshwar.dayal@hp.com

Sham Navathe
Georgia Institute of Technology
Atlanta, Georgia 30332, USA
E-mail: navathe@yahoo.com

Library of Congress Control Number: 2007931598

CR Subject Classification (1998): H.3.5, H.4.1, H.2.7, H.5.3, K.4.3, K.4.4, K.6, J.1

LNCS Sublibrary: SL 3 – Information Systems and Application, incl. Internet/Web and HCI

ISSN	0302-9743
ISBN-10	3-540-73949-1 Springer Berlin Heidelberg New York
ISBN-13	978-3-540-73949-4 Springer Berlin Heidelberg New York

This work is subject to copyright. All rights are reserved, whether the whole or part of the material is concerned, specifically the rights of translation, reprinting, re-use of illustrations, recitation, broadcasting, reproduction on microfilms or in any other way, and storage in data banks. Duplication of this publication or parts thereof is permitted only under the provisions of the German Copyright Law of September 9, 1965, in its current version, and permission for use must always be obtained from Springer. Violations are liable to prosecution under the German Copyright Law.

Springer is a part of Springer Science+Business Media
springer.com

© Springer-Verlag Berlin Heidelberg 2007
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 12100705 06/3180 5 4 3 2 1 0

Preface

The 1st Workshop on Business Intelligence for the Real-Time Enterprise (BIRTE 2006) was held on September 11, 2006 in conjunction with the 32nd International Conference on Very Large Data Bases (VLDB 2006) in Seoul, Korea. The co-location with VLDB is very important as the topic of the workshop was centered on different aspects in the lifecycle of business intelligence on very large enterprise-wide operational real-time data sets.

In today's competitive and highly dynamic environment, analyzing data to understand how the business is performing, to predict outcomes and trends, and to improve the effectiveness of business processes underlying business operations has become critical. The traditional approach to reporting is not longer adequate; users now demand easy-to-use intelligent platforms and applications capable of analyzing real-time business data to provide insight and actionable information at the right time. The end goal is to improve the enterprise performance by better and timelier decision making, enabled by the availability of up-to-date, high-quality information.

As a response, the notion of "real-time enterprise" has emerged and is beginning to be recognized in the industry. Gartner defines it as "using up-to-date information, getting rid of delays, and using speed for competitive advantage is what the real-time enterprise is all about... Indeed, the goal of the real-time enterprise is to act on events as they happen."

Although there has been progress in this direction and many companies are introducing products towards making this vision reality, there is still a long way to go. In particular, the whole lifecycle of business intelligence requires new techniques and methodologies capable of dealing with the new requirements imposed by the real-time enterprise. From the capturing of real-time business performance data to the injection of actionable information back into business processes, all the stages of the business intelligence (BI) cycle call for new algorithms and paradigms as the basis of new functionalities including dynamic integration of real-time data feeds from operational sources, evolution of ETL transformations and analytical models, and dynamic generation of adaptive real-time dashboards, just to name a few.

The goal of the BIRTE 2006 workshop was to provide a forum for the discussion of five major aspects of business intelligence for the real-time enterprise: Models and Concepts for Real-Time Enterprise Business Intelligence, Architectures for Real-Time Enterprise Business Intelligence, Uses Cases of Real-Time Enterprise Business Intelligence, Applications of Real-Time Enterprise Business Intelligence and Technologies for the Real-Time Enterprise Business Intelligence.

The workshop started with the keynote "Practical Considerations for Real-Time Business Intelligence" by Donovan Schneider. It continued with several sessions addressing various aspects of real-time data analysis. The first session "Streaming Data" concentrated on data streams as one mechanism for obtaining real-time enterprise data. The second session "Data Loading and Data Warehouse Architectures" addressed data loading and data warehouse architectures that both are a basis for the actual analysis task. The third session "Integration and Data Acquisition" focused on

heterogeneous data sources as well as mechanisms for obtaining real-time data. The final session “Business Processes and Contracts” extended the analysis aspect from data to processes. The workshop closed with the interesting panel “How Real Can Real-Time Business Intelligence Be?” moderated by Malu Castellanos, and Chi-Ming Chen, Mike Franklin, Minos Garofalakis, Wolfgang Lehner, Stuart Madnick and Krithi Ramamrithan as speakers.

The field of business intelligence for the real-time enterprise is fairly new, albeit increasingly important. This first workshop on the topic was meant to be a starting point of a series of several workshops covering various aspects in more detail over time. As academic research and industrial application experience more in-depth insights and use of this technology, an interesting research field opens up as well as an exciting area for practitioners. We encourage researchers and those in industry to continue their exciting work, and we encourage newcomers to enter this challenging and increasingly important field as there is still a lot of exciting work to be done.

We wish to express special thanks to the Program Committee members for providing their technical expertise in reviewing the submitted papers and preparing an interesting program. We are particularly grateful to the keynote speaker, Donovan Schneider, for delighting us with his very interesting keynote. Special recognition goes to the panelists for their enthusiastic participation in presenting their perspectives. To the authors of the accepted papers we express our appreciation for sharing their work and experiences in this workshop. Finally, we would like to extend many thanks to the VLDB 2006 Workshop Co-Chairs, Sang-goo Lee and Ming-Chien Shan, for their support in making this workshop possible.

September 2006

Christoph Bussler
Malu Castellanos
Umesh Dayal
Sham Navathe

Organization

Organizing Committee

General Chair

Umeshwar Dayal, Hewlett-Packard, USA

Program Committee Chairs

Christoph Bussler, Cisco Systems, Inc., USA

Malu Castellanos, Hewlett-Packard, USA

Sham Navathe, Georgia Institute of Technology, USA

Program Committee

Christof Bornhoevd, SAP Labs, USA

Mike Franklin, UC Berkeley, USA

Venkatesh Ganti, Microsoft, USA

Dimitrios Georgakopoulos, Telcordia Technologies, USA

Ramesh Jain, UC Irvine, USA

Meichun Hsu, HP Labs, China

Kamal Karlapalem, IIIT Hyderabad, India

Rajesh Parekh, Yahoo, USA

Torben B. Pedersen, Aalborg University, Denmark

Ee Peng, Nanyang Technological University, Singapore

Krithi Ramamritham, IIT Bombay, India

W.M.P. Van der Alst, Eindhoven University of Technology, The Netherlands

Panos Vassiliadis, University of Ioannina, Greece

Kazi Zaman, Siebel Systems Inc., USA

Publication Chair

Kamalakar Karlapalem, IIIT Hyderabad, India

Reviewers

Bin Zhang

Table of Contents

Practical Considerations for Real-Time Business Intelligence.....	1
<i>Donovan A. Schneider</i>	
What Can Hierarchies Do for Data Streams?	4
<i>Xuepeng Yin and Torben Bach Pedersen</i>	
Leveraging Distributed Publish/Subscribe Systems for Scalable Stream Query Processing	20
<i>Yongluan Zhou, Kian-Lee Tan, and Feng Yu</i>	
Transaction Reordering and Grouping for Continuous Data Loading	34
<i>Gang Luo, Jeffrey F. Naughton, Curt J. Ellmann, and Michael W. Watzke</i>	
A Scalable Heterogeneous Solution for Massive Data Collection and Database Loading	50
<i>Uri Shani, Aviad Sela, Alex Akilov, Inna Skarbovski, and David Berk</i>	
Two-Phase Data Warehouse Optimized for Data Mining	63
<i>Balázs Rácz, Csaba István Sidló, András Lukács, and András A. Benczúr</i>	
Document-Centric OLAP in the Schema-Chaos World	77
<i>Yannis Sismanis, Berthold Reinwald, and Hamid Pirahesh</i>	
Callisto: Mergers Without Pain	92
<i>Huong Morris, Hui Liao, Sriram Padmanabhan, Sriram Srinivasan, Eugene Kawamoto, Phay Lau, Jing Shan, and Ryan Wisnesky</i>	
Real-Time Acquisition of Buyer Behaviour Data – The Smart Shop Floor Scenario	106
<i>Bo Yuan, Maria Orlowska, and Shazia Sadiq</i>	
Business Process Learning for Real Time Enterprises	118
<i>Rodion Podorozhny, Anne Ngu, and Dimitrios Georgakopoulos</i>	
An Integrated Approach to Process-Driven Business Performance Monitoring and Analysis for Real-Time Enterprises.....	133
<i>Jonghun Park, Cheolkyu Jee, Kwanho Kim, Seung-Kyun Han, Duksoon Im, Wan Lee, and Noyoon Kim</i>	
Quality Contracts for Real-Time Enterprises	143
<i>Alexandros Labrinidis, Huiming Qu, and Jie Xu</i>	
Author Index	157