2000 International Database Engineering and Applications Symposium

Proceedings

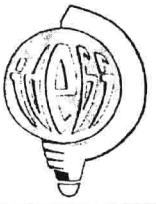

2000 International Database Engineering and Applications Symposium

September 18-20, 2000 • Yokohama, Japan

Sponsored by

Concordia University
Keio University
IDEAS

Edited by

Bipin C. Desai
Yasushi Kiyoki
Motomichi Toyama

Los Alamitos, California

Washington • Brussels • Tokyo

IEEE Computer Society Order Number PR00789
IEEE Order Plan Catalog Number PR00789
ISBN 0-7695-0789-1
Library of Congress Number 00-106721

Additional copies may be ordered from:

IEEE Computer Society
Customer Service Center
10662 Los Vaqueros Circle
P.O. Box 3014
Los Alamitos, CA 90720-1314
Tel: + 1 714 821 8380
Fax: + 1 714 821 4641
http://computer.org/
csbooks@computer.org

IEEE Service Center
445 Hoes Lane
P.O. Box 1331
Piscataway, NJ 08855-1331
Tel: + 1 732 981 0060
Fax: + 1 732 981 9667
http://shop.ieee.org/store/
customer-service@ieee.org

IEEE Computer Society
Asia/Pacific Office
Watanabe Bldg., 1-4-2
Minami-Aoyama
Minato-ku, Tokyo 107-0062
JAPAN
Tel: + 81 3 3408 3118
Fax: + 81 3 3408 3553
tokyo.ofc@computer.org

Editorial production by A. Denise Williams

Cover art production by Alex Torres

Printed in the United States of America by The Printing House, Inc.

Proceedings

2000 International Database Engineering and Applications Symposium

Foreword

IDEAS 2000 is the fourth in a series of meetings addressing the engineering aspect of databases and their applications for developing reliable systems for complex applications. The inaugural and the third IDEAS symposiums were held in Montreal, hosted by Concordia University. The second in this series was hosted by Cardiff University.

This volume contains the papers selected for presentation at this conference, hosted by Keio University. It includes 36 papers selected from some 90 submissions. In addition, it contains 12 short papers and the summaries of the invited talk and the keynote speech. Each of the submissions was reviewed by at least three members of the program committee, followed by on-line discussions. We would like to thank all authors who submitted their work to the symposium.

The database community is turning its attention to non-traditional applications and semi-structured data. This is reflected in the number of submissions dealing with the Web, XML, geographical, and other applications. These applications are pushing the frontier of knowledge and establishing novel ways of dealing with information overload as we enter the new millennium—the so-called "information age."

IDEAS has emerged as a valuable forum which addresses, in a timely manner, the complex problems of managing and mining data from databases. The engineering of new methodologies to provide rapid and reliable implementation and solutions to these problems involves drawing on the expertise in many areas of computer science, including the growing trend in engineering of reliable and safe software systems. It is gratifying to see that the database research community has adopted the IDEAS symposia to address current issues related to the information age by reporting on the experience in building reliable advanced database applications, user interfaces to these applications, emerging data modeling and database theories, and technology.

I wish to express my gratitude to the many people at Keio University, particularly Professors Kiyoki and Toyama, who have provided active support for IDEAS 2000, and the local organizing committee, for arranging this meeting on short notice and providing excellent facilities for it. The Program Committee and the referees have ensured the quality of the symposium through meticulous referring of the submitted papers. I would also like to thank Concordia University, Keio University, and IEEE Computer Society for providing support to make IDEAS 2000 possible.

Bipin C. Desai

General Chair

Conference Organizers

General Chair and Program Co-Chair

Bipin C. Desai
Concordia University

Local Organizing Chair and Program Co-Chair

Yasushi Kiyoki
Keio University

Publicity Chair

Teruhisa Hochin
Fukui University

Local Arrangements Chair

Motomichi Toyama,
Keio University

IDEAS 2000 Reviewers

Program Committee

Michel Adiba, *Grenoble Univ.-IMAG-LSR*
Alex Biliris, *AT&T Research*
Joachim Biskup, *Dortmund*
Christine Collet, *LSR-IMAG*
Isabel Cruz, *WPI*
Bipin C. Desai, *Concordia*
Anne Doucet, *Pierre et Marie Curie Paris-6*
Prabhakar Gokul, *Lucent*
Gosta Grahne, *Concordia*
William I Grosky, *Wayne State Univ.*
Jiawei Han, *Simon Fraser Univ*
Teruhisa Hochin, *Fukui Univ.*
Yoshiharu Ishikawa, *Tsukuba Univ.*
Bala Iyer, *IBM Santa Teresa Lab*
Hasan Jamil, *Mississippi State*
Leonid A. Kalinichenko, *Russian Acad Sci*
Hiroyuki Kawano, *Kyoto Univ.*
Brigitte Kerherve, *UQAM*
Yasushi Kiyoki, *Keio Univ.*
Laks V. S. Lakshmanan, *Concordia*
Georg Lausen, *Freiburg*
Tok Wang Ling, *NUS-Singapore*
Mengchi Liu, *Regina*
Sanjay K Madria, *NTU-Singapore*
Richard McClatchey, *CERN*
Mukesh Mohania, *Western Michigan*
Sylvia Osborn, *Western Ont.*
Jaroslav Pokorny, *Charles Univ.*
Zbigniew W Ras, *UNCC*
Louiqa Raschid, *Maryland*
Peter Revesz, *Univ. Nebraska-Lincoln*
Domenico Sacca, *Calabria*
Fereidoon Sadri, *UNCG*
Timos Sellis, *NTU:Athens*
Divesh Srivastava, *AT&T Research*
Kazimierz Subieta, *Polish Acad Sci.*
Kazutoshi Sumiya, *Kobe Univ.*
Roger Tagg, *Massey-NZ*
Riccardo Torlone, *Universitá di Roma Tre*
Motomichi Toyama, *Keio Univ.*
Nectaria Tryfona, *Aalborg University*
Can Tuerker, *ETH-Zurich*
John Velissarios, *Price Waterhouse*
Gottfried Vossen, *Univ. of Muenster*
Vlad Wietrzyk, *Macquarie-Australia*

External Reviewers

Sihem Amer-Yahia, *AT&T Labs*
Hong Chen, *Regina*
Marlon Dumas, *Grenoble*
Pierre Habraken, *Grenoble*
Flip Korn, *AT&T Labs*
Hervé Martin, *Grenoble*
Hongwei Qi, *Regina*
Dieter Pfoser, *Aalborg University*
Simonas Saltenis, *Aalborg University*
Yibin Su, *Regina*
Keishi Tajima, *Kobe University*
José-Luis Zechinelli, *Grenoble*
Zhong Zhang, *Simon Fraser*

Table of Contents

2000 International Database Engineering and Applications Symposium

Session 1

Opening Talk

Business-to-Business Connectivity and Emerging Protocols for E-Biz

Daniel Dias
IBM T. J. Watson Research Center
Yorktown Hts, NY
email: dias@us.ibm.com

Summary

For many years EDI served as the dominant protocol for e-biz over private networks. Due to the rapid spread of the Internet, several new Business-to-Business (B2B) interaction protocols have become popular, e.g., Open Buying on the Internet (OBI), RosettaNet, and Vendor-defined protocols such as Ariba's Punchout process. We will describe these protocols, with an emphasis on their similarity and differences. B2B marketplaces are emerging as hubs through which suppliers connect to buyers. We outline the connectivity requirements between these marketplaces, and the Marketplace-to-Marketplace (M2M) protocols that are required.

We show how all of these emerging protocols can be specified in XML using a trading partner agreement (TPA). The TPA specifies three layers in the B2B protocol: The transport, document exchange and business protocol layers. Essentially, the transport layer provides a choice of transport such as HTTP, SMTP, FTP, and specifies related information; the document exchange layer handles different message formats, transformations, application level security including authentication and non-repudiation, and provides services such as audit trails and document repositories; the business protocol layer specifies the sequence of the B2B interactions, responsiveness criteria, and error handling.

From the TPA, an instantiation of the B2B protocol can be built. We describe a Business to Business Protocol Framework (BPF), on which different protocols across autonomous businesses can be instantiated. Our work has the potential to decrease the time of implementation of e-biz solutions. BPF provides a comprehensive set of tools and enablers for ease of specification, configuration, plug-in, customization, and execution of a set of TPAs between business partners. It is the gateway and coordinator across intra-business and inter-business processes, (e.g., between the buy/sell component of a local business, the remote businesses, and the back-end systems). The protocols are expressed as an electronic TPA in XML using a higher level authoring tool and are registered to the BPF server along with the internal business processes for setting up such B-B interactions. BPF uses database services for various functions including guaranteed delivery, recovery and message repositories. As examples, we show how OBI and Ariba's punchout processes can be specified with TPAs and implemented with BPF.

0-7695-0789-1/00 $10.00 © 2000 IEEE

Session 2

Formalism

A Formal Background to build Constraint Objects

Annalisa Di Deo, Dmitri Boulanger
GMD-First
Berlin, Germany
{anna, dmitri}@first.gmd.de

Abstract

Constraint databases exploit a fundamental duality: a constraint (first-order formula) with free variables is interpreted as a set of database tuples that satisfy it. And, conversely, a database object can be viewed as a constraint. This enables us to use constraints as basic data types in the underlying DBMS and to enjoy expressiveness of first-order logic and most advantages from constraint programming.

Most of the existing frameworks consider the relational data model as a basic data representation model and use linear constraints to represent complex objects. This is not always efficient. Moreover, modern DBMS offer much more than a simple set of relational tables. Therefore, we focus on the definition and formal background for an implementation of a constraint data model to be built on the top of an object-relational DBMS. As an illustration, our approach is capable to exploit existing spatial tools (like the one from Oracle8) yielding the declarativeness and expressiveness of constraint programming.

Keywords: *constraint databases, cylindric algebras, constraint objects, spatial relations.*

1. Introduction

With the increasing complexity of applications that store and manage spatial and temporal data, Database Management Systems are facing new challenges to efficiently represent and query such data, preserving the declarative and user-friendly features of the query languages. It has been acknowledged for a long time that the relational model fails to handle highly complex data and their operations. For this reason, the existing commercial DBMS are addressing the use of object-oriented models, and provide tools to represent complex objects and to reason with them efficiently. For instance, in the area of spatial applications, Oracle8 has developed a Spatial Cartridge ([1]) that no longer manages complex geometrical shapes as sets of points stored in the tables, but uses specially designed data types and operations to handle spatial data efficiently.

Spatial data constitute an example of multidimensional data to which a large effort has been devoted. In most cases [15, 18] the models consist in extending relational DBMSs with abstract spatial data types encapsulating geometric structures and operations. Due to the extensive range of applications, the problem of developing and unifying data models still remains open, when a facility to manipulate different data in a uniform way is needed.

The recent fields of constraint databases, initiated at the beginning of the decade [16], has lead to sound data models and query languages for multi-dimensional data [10, 12, 15]. The big challenge of constraint databases is to introduce constraints as basic data type in databases. The main principles are based on a simple and fundamental duality that a constraint can be seen as a relation, or, vice versa, that a relation can be interpreted as a simple form of constraint. This enables us to manipulate relations in the underlying DMBS in a symbolic way enjoying expressive power of first-order logic. That is, the syntax consists of constraints, i.e. symbolic expressions; the semantics are the corresponding relations. From this data modeling perspective, constraints are used as a unifying data type for the representation of different sorts of multi-dimensional data.

Most of the existing frameworks rely on linear constraints, first appeared in [13, 20]. They allow to manipulate relations of arbitrary dimension between points in a symbolic manner. In spite of their expressiveness, the existing constraint databases are not always efficient enough for real problems where the size and the complexity of data are high. In particular, spatial data might be so complex, e.g. non-convex geometries with multiple holes, that constraint database models cannot handle them efficiently and need to refer to optimized structures of the objects and operations

0-7695-0789-1/00 $10.00 © 2000 IEEE