

Lecture Notes in Computer Science

2572

Edited by G. Goos, J. Hartmanis, and J. van Leeuwen

Springer

Berlin

Heidelberg

New York

Barcelona

Hong Kong

London

Milan

Paris

Tokyo

Diego Calvanese Maurizio Lenzerini
Rajeev Motwani (Eds.)

Database Theory – ICDT 2003

9th International Conference
Siena, Italy, January 8-10, 2003
Proceedings



Springer

Series Editors

Gerhard Goos, Karlsruhe University, Germany
Juris Hartmanis, Cornell University, NY, USA
Jan van Leeuwen, Utrecht University, The Netherlands

Volume Editors

Diego Calvanese
Maurizio Lenzerini
Dip. di Informatica e Sistemistica
Università di Roma "La Sapienza"
Via Salaria 113, 00198 Roma, Italy
E-mail: {calvanese/lenzerini}@dis.uniroma1.it

Rajeev Motwani
Department of Computer Science, Room 474
Gates Computer Science Building 4B
Stanford University, Stanford, CA 94305-9045, USA
E-mail: rajeev@cs.stanford.edu

Cataloging-in-Publication Data applied for

A catalog record for this book is available from the Library of Congress.

Bibliographic information published by Die Deutsche Bibliothek
Die Deutsche Bibliothek lists this publication in the Deutsche Nationalbibliografie;
detailed bibliographic data is available in the Internet at <<http://dnb.ddb.de>>.

CR Subject Classification (1998): H.2, F.1.3, F.4.1, I.2.1, H.4, F.2, H.3

ISSN 0302-9743

ISBN 3-540-00323-1 Springer-Verlag 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-Verlag. Violations are liable for prosecution under the German Copyright Law.

Springer-Verlag Berlin Heidelberg New York
a member of BertelsmannSpringer Science+Business Media GmbH

<http://www.springer.de>

© Springer-Verlag Berlin Heidelberg 2003
Printed in Germany

Typesetting: Camera-ready by author, data conversion by PTP Berlin, Stefan Sossna e. K.
Printed on acid-free paper SPIN: 10872297 06/3142 5 4 3 2 1 0

Preface

This volume collects the papers presented at ICDT 2003, the 9th International Conference on Database Theory, held from 8 to 10 January 2003 in the Rettorato dell'Università di Siena, Siena, Italy.

ICDT (<http://alpha.luc.ac.be/~lucp1080/icdt/>) has now a long tradition of international conferences, providing a biennial scientific forum for the communication of high-quality research results on theoretical aspects of all forms of database systems and database technology. ICDT is traditionally held in historic European locations: Rome in 1986, Bruges in 1988, Paris in 1990, Berlin in 1992, Prague in 1995, Delphi in 1997, Jerusalem in 1999, and London in 2001. ICDT has merged with the Symposium on Mathematical Fundamentals of Database Systems (MFDBS), initiated in Dresden in 1987, and continued in Visegrad in 1989 and Rostock in 1991.

This volume contains 26 papers selected from 92 submissions, plus three invited papers by Hector Garcia-Molina, Yannis Ioannidis, and Limsoon Wong. We wish to thank all the authors who submitted papers, the members of the program committee for their efforts in reviewing and selecting the papers, the external referees, the organizing committee, and the sponsors for their support.

Diego Calvanese
Maurizio Lenzerini
Rajeev Motwani

Organization

ICDT 2003 was organized by the Dipartimento di Informatica e Sistemistica, Università di Roma “La Sapienza”.

Organizing Committee

Maurizio Lenzerini (Chair)
Diego Calvanese
Marco Gori (Local Chair)
Tiziana Toni

Program Co-chairs

Maurizio Lenzerini (University of Rome “La Sapienza”, Italy)
Rajeev Motwani (Stanford University, U.S.A.)

Program Committee

Serge Abiteboul (INRIA, France)
Foto Afrati (National Technical University of Athens, Greece)
Catriel Beeri (Hebrew University of Jerusalem, Israel)
Stephan Bressan (National University of Singapore, Singapore)
Peter Buneman (Univ. of Edinburgh, UK, and Univ. of Pennsylvania, USA)
Diego Calvanese (University of Rome “La Sapienza”, Italy)
Surajit Chaudhuri (Microsoft Research)
Ronald Fagin (IBM Almaden Research Center, USA)
Alon Y. Halevy (University of Washington, USA)
Heikki Mannila (University of Helsinki, Finland)
Alberto O. Mendelzon (University of Toronto, Canada)
Guido Moerkotte (University of Mannheim, Germany)
Jeffrey F. Naughton (University of Wisconsin, USA)
Yannis Papakonstantinou (University of California, USA)
Luigi Palopoli (University of Reggio Calabria, Italy)
Domenico Saccà (University of Calabria, Italy)
Vladimir Sazonov (University of Liverpool, UK)
Bernhard Thalheim (Brandenburg Technical University at Cottbus, Germany)
Jeffrey D. Ullman (Stanford University, USA)
Jan Van den Bussche (Limburgs Universitair Centrum, Belgium)
Vasilis Vassalos (New York University, USA)

External Referees

Fabrizio Angiulli
Marcelo Arenas
Brian Babcock
Stefano Basta
Michael Benedikt
Angela Bonifati
Francesco Buccafurri
Marco Cadoli
Andrea Calì
Paolo Ciaccia
Gregory Cobena
Gautam Das
Mayur Datar
Giuseppe De Giacomo
Wenfei Fan
Sergio Flesca
Paolo Franciosa
Irina Fundulaki
Filippo Furfaro
Venkatesh Ganti
Fosca Giannotti
Aris Gionis
Jonathan Goldstein
Dimitrios Gunopulos
Gianluigi Greco
Sergio Greco
Antonella Guzzo
Sven Helmer
Christoph Koch
Phokion Kolaitis

Nick Koudas
Gabriel Kuper
Giovambattista Ianni
Pasquale Legato
Domenico Lembo
Nicola Leone
Chen Li
Leonid Libkin
Alexei Lisitsa
Giuseppe Manco
Gerome Miklau
Tova Milo
Vivek Narasayya
Clara Pizzuti
Luigi Pontieri
Andrea Pugliese
Maurizio Rafanelli
Domenico Rosaci
Riccardo Rosati
Massimo Ruffolo
Pasquale Rullo
Alan P. Sexton
Francesco Scarcello
Luc Segoufin
Jerome Simeon
Domenico Talia
Domenico Ursino
Suresh Venkatasubramanian
Victor Vianu
Ester Zumpano

Table of Contents

Invited Papers

Open Problems in Data-Sharing Peer-to-Peer Systems	1
<i>Neil Daswani, Hector Garcia-Molina, Beverly Yang</i>	
Approximations in Database Systems	16
<i>Yannis Ioannidis</i>	
Bioinformatics Adventures in Database Research	31
<i>Jinyan Li, See-Kiong Ng, Limsoon Wong</i>	

Reasoning about XML Schemas and Queries

Incremental Validation of XML Documents	47
<i>Yannis Papakonstantinou, Victor Vianu</i>	
Typechecking Top-Down Uniform Unranked Tree Transducers	64
<i>Wim Martens, Frank Neven</i>	
Structural Properties of XPath Fragments	79
<i>Michael Benedikt, Wenfei Fan, Gabriel M. Kuper</i>	
On Reasoning about Structural Equality in XML: A Description Logic Approach	96
<i>David Toman, Grant Weddell</i>	

Aggregate Queries

Containment of Aggregate Queries	111
<i>Sara Cohen, Werner Nutt, Yehoshua Sagiv</i>	
Auditing Sum Queries	126
<i>Francesco M. Malvestuto, Mauro Mezzini</i>	
CRB-Tree: An Efficient Indexing Scheme for Range-Aggregate Queries	143
<i>Sathish Govindarajan, Pankaj K. Agarwal, Lars Arge</i>	
Optimal Range Max Datacube for Fixed Dimensions	158
<i>Chung Keung Poon</i>	

Query Evaluation

Processing XML Streams with Deterministic Automata	173
<i>Todd J. Green, Gerome Miklau, Makoto Onizuka, Dan Suciu</i>	

Deciding Termination of Query Evaluation in Transitive-Closure
 Logics for Constraint Databases 190
Floris Geerts, Bart Kuijpers

Query Rewriting and Reformulation

Data Exchange: Semantics and Query Answering 207
*Ronald Fagin, Phokion G. Kolaitis, Renée J. Miller,
 Lucian Popa*

Reformulation of XML Queries and Constraints 225
Alin Deutsch, Val Tannen

New Rewritings and Optimizations for Regular Path Queries 242
Gösta Grahne, Alex Thomo

Database Interrogation Using Conjunctive Queries 259
Michał Bielecki, Jan Van den Bussche

Semistructured versus Structured Data

On the Difficulty of Finding Optimal Relational Decompositions
 for XML Workloads: A Complexity Theoretic Perspective 270
*Rajasekar Krishnamurthy, Venkatesan T. Chakaravarthy,
 Jeffrey F. Naughton*

Generating Relations from XML Documents 285
Sara Cohen, Yaron Kanza, Yehoshua Sagiv

Query Containment

Containment for XPath Fragments under DTD Constraints 300
Peter T. Wood

XPath Containment in the Presence of Disjunction, DTDs,
 and Variables 315
Frank Neven, Thomas Schwentick

Decidable Containment of Recursive Queries 330
Diego Calvanese, Giuseppe De Giacomo, Moshe Y. Vardi

Containment of Conjunctive Queries with Safe Negation 346
Fang Wei, Georg Lausen

Consistency and Incompleteness

Probabilistic Interval XML 361
Edward Hung, Lise Getoor, V.S. Subrahmanian

Condensed Representation of Database Repairs for Consistent Query Answering	378
<i>Jef Wijsen</i>	
Typing Graph-Manipulation Operations	394
<i>Jan Hidders</i>	
Characterizing the Temporal and Semantic Coherency of Broadcast-Based Data Dissemination	410
<i>Evaggelia Pitoura, Panos K. Chrysanthis, Krithi Ramamritham</i>	
Data Structures	
An Efficient Indexing Scheme for Multi-dimensional Moving Objects	425
<i>Khaled Elbassioni, Amr Elmasry, Ibrahim Kamel</i>	
Nearest Neighbors Can Be Found Efficiently If the Dimension Is Small Relative to the Input Size	440
<i>Michiel Hagedoorn</i>	
Author Index	455