

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

New York University, NY, 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

Wolfgang Lindner Marco Mesiti
Can Türker Yannis Tzitzikas
Athena Vakali (Eds.)

Current Trends in Database Technology – EDBT 2004 Workshops

**EDBT 2004 Workshops PhD, DataX, PIM, P2P&DB, and ClustWeb
Heraklion, Crete, Greece, March 14-18, 2004
Revised Selected Papers**



Springer

Volume Editors

Wolfgang Lindner
Massachusetts Institute of Technology, CSAIL
32 Vassar Street, Cambridge, MA 02139, USA
E-mail: wolfgang@csail.mit.edu

Marco Mesiti
University of Milan, DICO
Via Comelico 39/41, 20135 Milan, Italy
E-mail: mesiti@ dico.unimi.it

Can Türker
Swiss Federal Institute of Technology Zurich
8092 Zurich, Switzerland
E-mail: tuerker@inf.ethz.ch

Yannis Tzitzikas
University of Namur, Institute of Informatics
Rue Grandgagnage 21, 5000 Namur, Belgium
E-mail: ytz@info.fundp.ac.be

Athena Vakali
Aristotle University, Department of Informatics
54124 Thessaloniki, Greece
E-mail: avakali@csd.auth.gr

Library of Congress Control Number: 2004115768

CR Subject Classification (1998): H.2, H.3, H.4, H.5.1, C.2.4, K.4.4

ISSN 0302-9743
ISBN 3-540-23305-9 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
springeronline.com

© Springer-Verlag Berlin Heidelberg 2004
Printed in Germany

Typesetting: Camera-ready by author, data conversion by Scientific Publishing Services, Chennai, India
Printed on acid-free paper SPIN: 11328957 06/3142 5 4 3 2 1 0

Preface

This volume comprises papers from the following five workshops that were part of the complete program for the International Conference on Extending Database Technology (EDBT) held in Heraklion, Greece, March 2004:

- ICDE/EDBT Joint Ph.D. Workshop (PhD)
- Database Technologies for Handling XML-information on the Web (DataX)
- Pervasive Information Management (PIM)
- Peer-to-Peer Computing and Databases (P2P&DB)
- Clustering Information Over the Web (ClustWeb)

Together, the five workshops featured 61 high-quality papers selected from approximately 180 submissions. It was, therefore, difficult to decide on the papers that were to be accepted for presentation. We believe that the accepted papers substantially contribute to their particular fields of research. The workshops were an excellent basis for intense and highly fruitful discussions. The quality and quantity of papers show that the areas of interest for the workshops are highly active. A large number of excellent researchers are working on the aforementioned fields producing research output that is not only of interest for other researchers but also for industry. The organizers and participants of the workshops were highly satisfied with the output. The high quality of the presenters and workshop participants contributed to the success of each workshop. The amazing environment of Heraklion and the location of the EDBT conference also contributed to the overall success. Last, but not least, our sincere thanks to the conference organizers – the organizing team was always willing to help and if there were things that did not work, assistance was quickly available.

On the basis of the reviews and discussions during the workshops, 55 authors were asked to submit revised versions of their workshop papers for these LNCS postproceedings. When revising their papers, authors were asked to consider not only the comments of the reviewers but also the comments of other workshop participants. The revised papers underwent a second reviewing process before they were accepted for these postproceedings.

ICDE/EDBT Ph.D. Workshop (PhD)

Continuing in its tradition, the Ph.D. Workshop brought together Ph.D. students in the field of database technology outside of the ICDE/EDBT conference series. It offers Ph.D. students the opportunity to present, discuss, and receive feedback on their research in a constructive and international environment.

For the first time the Ph.D. Workshop was actually a real distributed event. The first session was held in conjunction with the 9th International Conference on Extending Database Technology (EDBT) on March 18, 2004, in Heraklion (Greece). The second

session was held in conjunction with the 20th International Conference on Data Engineering (ICDE) on March 29, 2004, in Boston (USA).

As for the statistics: we received 86 submissions (from 21 countries). After a careful review process (each paper was evaluated by three committee members followed by an online discussion), 21 papers were selected for presentation at the workshop, out of which 18 contributors were invited to submit revised and extended versions of their papers for these proceedings.

Workshop Chair

Wolfgang Lindner

Massachusetts Institute of Technology (USA)

Program Committee

Karl Aberer	EPFL Lausanne (Switzerland)
Anastassia Ailamaki	Carnegie Mellon University (USA)
Walid G. Aref	Purdue University (USA)
Phil Bernstein	Microsoft (USA)
Elisa Bertino	Purdue University (USA)
Michael Böhlen	Free University of Bolzano-Bozen (Italy)
Stefano Ceri	University of Milan (Italy)
Klaus Dittrich	University of Zurich (Switzerland)
Georgios Evangelidis	University of Macedonia (Greece)
Michael Franklin	UC Berkeley (USA)
Dina Goldin	University of Connecticut (USA)
Christian S. Jensen	Aalborg University (Denmark)
Björn Þór Jónsson	Reykjavik University (Iceland)
Martin Kersten	CWI (Netherlands)
Tan Kian Lee	NUS (Singapore)
Peter Lockemann	University of Karlsruhe (Germany)
Bertram Ludaescher	UC San Diego (USA)
Samuel Madden	MIT (USA)
David Maier	Oregon Health and Science University (USA)
Ioana Manolescu	INRIA (France)
Holger Meyer	University of Rostock (Germany)
Klaus Meyer-Wegener	University of Erlangen-Nuremberg (Germany)
Felix Naumann	Humboldt University Berlin (Germany)
Erich Neuhold	IPSI Fraunhofer (Germany)
Dimitris Papadias	Hong Kong University (China)
Jan Paredaens	University of Antwerp (Belgium)
Neoklis Polyzotis	UC Santa Cruz (USA)
Ken Salem	University of Waterloo (Canada)
Timos Sellis	National Technical University of Athens (Greece)
Heiko Schuld	UMIT (Austria)
Peter Triantafillou	University of Patras (Greece)
David Toman	University of Waterloo (Canada)

Database Technologies for Handling XML-Information on the Web (DataX)

Since XML was proposed by the W3C in 1998, the database community has been working on ways to manage semistructured information by extending traditional database systems and by proposing new native XML-based systems in order to store, maintain, exchange, and securely access XML documents. DataX brought together experts from several fields of information technology to discuss new interesting results and applications of XML data management. An invited paper reports the current assessment of the area and outlines the promising challenges for the next few years. Moreover, the technical program addresses important topics concerning querying and indexing of XML sources along with the evolution of XML schema and applications. Finally, a panel deals with important questions in XML data management.

Workshop Chairs

Barbara Catania	University of Genoa (Italy)
Akmal B. Chaudhri	IBM developerWorks (UK)
Giovanna Guerrini	University of Pisa (Italy)
Marco Mesiti	University of Milan (Italy)

Program Committee

Ricardo Baeza-Yates	University of Chile (Chile)
Zohra Bellahsene	LIRMM (France)
Stéphane Bressan	NUS (Singapore)
Barbara Catania	University of Genoa (Italy)
Akmal B. Chaudhri	IBM developerWorks (UK)
Klaus R. Dittrich	University of Zurich (Switzerland)
Elena Ferrari	University of Insubria (Italy)
Mary F. Fernandez	AT&T Research (USA)
Giorgio Ghelli	University of Pisa (Italy)
Sergio Greco	University of Calabria (Italy)
Giovanna Guerrini	University of Pisa (Italy)
Hosagrahar V. Jagadish	University of Michigan (USA)
Zoé Lacroix	Arizona State University (USA)
Anna Maddalena	University of Genoa (Italy)
Gianni Mecca	University of Basilicata (Italy)
Marco Mesiti	University of Milan (Italy)
M. Tamer Özsu	University of Alberta (Canada)
Rajeev Rastogi	Bell Labs (USA)
Ralf Schenkel	Max-Planck Institute (Germany)
Divesh Srivastava	AT&T Research (USA)
Rainer Unland	University of Essen (Germany)
Athena Vakali	Aristotle University of Thessaloniki (Greece)

External Reviewers

Bernd Amann	Elio Masciari	Cristina Sirangelo
Giovanni Conforti	Pietro Mazzoleni	Dominique Stein
Sergio Flesca	Andrea Perego	Andrea Tagarelli
Ela Hunt	Mark Roantree	Pierangelo Veltri
Anne Laurent	Awny Sayed	Alessandro Verri
Paolo Manghi	Carlo Sartiani	Ning Zhang

Pervasive Information Management (PIM)

Wireless connectivity, along with increasingly small and powerful mobile devices and sensors, enables a wide range of new applications that will radically change the way information is managed and processed today. Information becomes ubiquitous, highly distributed and at the same time accessible from everywhere at any time. Information access takes place in highly dynamic and unstable networks. Nevertheless, users and application developers will expect information processing to continue under similar guarantees as those offered by today's stationary and more or less centralized systems, even if some nodes of the pervasive information network are (temporarily) disconnected and/or are in motion. Examples of these guarantees are those given by database management systems, e.g., consistency and durability of data. Additional challenges stem from the fact that mobile devices, in particular embedded devices and sensors, are less powerful than their stationary counterparts; they are smaller (resulting, for example, in smaller in- and output devices and less available storage), have restricted energy supplies, and communicate via expensive and unreliable wireless communication media.

The papers accepted for this postconference proceedings address a broad range of issues in pervasive information systems. New and existing concepts and techniques are developed in the light of the rapidly increasing mobility of users and the great advances in system infrastructures, mobile devices, and sensor technologies.

Workshop Chairs

Can Türker	ETH Zurich (Switzerland)
Birgitta König-Ries	TU München (Germany)

Program Committee

Karl Aberer	EPFL Lausanne (Switzerland)
Michel E. Adiba	University of Grenoble (France)
Christoph Gollmick	University of Jena (Germany)
Hagen Höpfer	University of Magdeburg (Germany)
Valérie Issarny	INRIA Rocquencourt (France)

Christian S. Jensen	Aalborg University (Denmark)
Birgitta König-Ries	TU München (Germany)
Sanjay Kumar Madria	University of Missouri-Rolla (USA)
Rainer Malaka	EML Heidelberg (Germany)
Kia Makki	Florida International University (USA)
Eduardo Mena	University of Zaragoza (Spain)
Evaggelia Pitoura	University of Ioannina (Greece)
Kurt Rothermel	University of Stuttgart (Germany)
George Samaras	University of Cyprus (Cyprus)
Kai-Uwe Sattler	University of Ilmenau (Germany)
Heiko Scholdt	UNIT Innsbruck (Austria)
Can Türker	ETH Zurich (Switzerland)
Özgür Ulusoy	Bilkent University (Turkey)
Jari Veijalainen	University of Jyväskylä (Finland)
Ouri Wolfson	University of Illinois (USA)

External Reviewers

Hidir Aras	Klaus Haller	Gregor Schiele
Martin Bauer	Martin Heusse	Roman Schmidt
Hu Cao	Sergio Ilarri	Chara Skouteli
Christos Christophi	Juhong Liu	Jie Wu
Tobias Drosdol	Philipp Obreiter	Bo Xu
Dimosthenis Georgiadis	Christoforos Panayiotou	Huabei Yin
Tirthankar Ghosh	Robert Porzel	

Peer-to-Peer Computing and Databases (P2P&DB)

The P2P paradigm holds many promises: it couples naturally with the Internet, the universal knowledge and service exchange medium; it favors scalability, by allowing the seamless plugging of data, services and computational resources into the global system; it increases system resilience, by avoiding unique points of failures; and it can speed up global access by distributing the indexing and query processing tasks to multiple computing nodes. This is all in contrast to the centralized architecture of current systems offering global access, such as search engines, Web directories, or mediators. The potential for progress is very encouraging, and there are already several systems that subscribe to the P2P computing paradigm, to some degree. On a theoretical level, researchers have started looking into models and algorithms to tackle basic problems in the P2P framework, as witnessed by the increasing number of papers in international workshops and conferences. This workshop, co-located with EDBT 2004 and focused

on peer-to-peer systems and database computing, was no exception. It attracted quality papers, most of which provoked lively discussions amongst the participants.

Workshop Chairs

Carlo Meghini	ISTI-CNR (Italy)
Nicolas Spyrtos	University of Paris-Sud (France)
Yannis Tzitzikas	University of Namur (Belgium)

Program Committee

Karl Aberer	EPFL Lausanne (Switzerland)
Anastasia Analyti	ICS-FORTH (Greece)
Paolo Atzeni	University of Rome 3 (Italy)
Diego Calvanese	Free University of Bozen-Bolzano (Italy)
Vassilis Christophides	ICS-FORTH (Greece)
Norbert Fuhr	University of Duisburg (Germany)
Mohand-Said Hacid	University Claude Bernard Lyon 1 (France)
Hatem Haddad	VTT (Finland)
Klaus P. Jantke	DFKI GmbH (Germany)
Ross King	Austrian Research Centers (Austria)
Matthias Klusch	DFKI GmbH (Germany)
Manolis Koubarakis	Technical University of Crete (Greece)
Heikki Mannila	Helsinki University of Technology (Finland)
John Mylopoulos	University of Toronto (Canada)
Wolfgang Nejdl	University of Hannover (Germany)
Evaggelia Pitoura	University of Ioannina (Greece)
Philippe Rigaux	University of Paris-Sud (France)
Thomas Risse	IPSI Fraunhofer (Germany)
Marie-Christine Rousset	University of Paris-Sud (France)
Luciano Serafini	Cultural Institute of Trentino (Italy)
Vagan Terziyan	University of Jyväskylä (Finland)

External Reviewers

Josenildo Costa da Silva	Predrag Knežević	Henrik Nottelmann
Philippe Cudré-Mauroux	Georgia Koloniari	Loris Penserini
Sarunas Girdzijauskas	Jafar Movahedi	Christos Tryfonopoulos

Clustering Information over the Web (ClustWeb)

Web data management has become a critical emerging research area, due to the exponential increase in the information circulation and dissemination over the Web. Web data mining is an evolving field of high interest to a wide academic and technical community.

The ClustWeb workshop focused on the topic of clustering information over the Web, as an emerging topic in current Web data mining efforts. The papers presented and the panel presentation contributed to understanding the role of clustering mechanisms and methodologies in accessing Web information (such as documents, link paths, and users on the Web). The topics of the workshop are summarized in:

- clustering structured Web sources,
- similarity ranking on the Web graph,
- clustering query logs (in search engines), towards dissemination of information,
- protecting Web data in the framework of clustering

The ClustWeb workshop addressed the issues involved in the effect of Web data clustering on increasing Web information accessibility, decreasing lengths in Web navigation pathways, improving Web user request servicing, integrating various data representation standards, and extending current Web information organization practices.

Workshop Chairs

Jaroslav Pokorný
Athena Vakali

Charles University of Prague (Czech Republic)
Aristotle University of Thessaloniki (Greece)

Program Committee

Eleftherios Angelis	Aristotle University of Thessaloniki (Greece)
Boualem Benatallah	University of New South Wales in Sydney (Australia)
Salima Benbernou	University of Lyon 1 (France)
Igor Cadez	University of California, Irvine (USA)
Barbara Catania	University of Genoa (Italy)
Siobhán Clarke	Trinity College, Dublin (Ireland)
Elena Ferrari	University of Insubria (Italy)
Mohand-Said Hacid	University of Lyon 1 (France)
Carlos Hurtado	University of Chile (Chile)
Nick Koudas	AT&T Labs (USA)
Fang Li	Shanghai Jiao Tong University (China)
András Lukács	Hungarian Academy of Sciences (Hungary)
Marco Mesiti	University of Milan (Italy)
Apostolos Papadopoulos	Aristotle University of Thessaloniki (Greece)
Jaroslav Pokorný	Charles University (Czech Republic)
Yücel Saygin	Sabancı University (Turkey)
Václav Snášel	Technical University of Ostrava (Czech Republic)
Athena Vakali	Aristotle University of Thessaloniki (Greece)

Acknowledgments

We would sincerely like to thank all program committee members, as well as all external referees for their excellent work in evaluating the submitted papers. Also, we would like to thank all authors for submitting a paper to one of our workshops. Special thanks go to Andrea Perego from the University of Milan (Italy) who did an outstanding job in compiling these proceedings, as well as to Alfred Hofmann from Springer for his cooperation and advice in putting this volume together.

Cambridge, USA
July 2004

Wolfgang Lindner (Program Chair PhD)
Marco Mesiti (Program Co-chair DataX)
Can Türker (Program Chair PIM)
Yannis Tzitzikas (Program Co-chair P2P&DB)
Athena Vakali (Program Co-chair ClustWeb)

Table of Contents

ICDE/EDBT Joint PhD Workshop (PhD)

Session I: March 18, 2004, Heraklion, Crete, Greece

Advanced Query Processing and Optimization

Querying Sliding Windows Over Online Data Streams <i>Lukasz Golab</i>	1
MIRA: Multilingual Information Processing on Relational Architecture <i>A. Kumaran</i>	12

Semi-structured Data and XML

Index-Based Keyword Search in Mediator Systems <i>Ingolf Geist</i>	24
Concept-Based Search on Semi-structured Data Exploiting Mined Semantic Relations <i>Jens Graupmann</i>	34

Data Model, Access Methods, and Control

Distributed and Scalable Similarity Searching in Metric Spaces <i>Michal Batko</i>	44
Pattern Based Management: Data Models and Architectural Aspects <i>Anna Maddalena</i>	54

Semantic Web

Managing Dynamic Repositories for Digital Content Components <i>André Santanchè, Claudia Bauzer Medeiros</i>	66
Semantic Web Recommender Systems <i>Cai-Nicolas Ziegler</i>	78

Session II: March 29, 2004, Boston, MA, USA

Information Hiding and Trust Negotiation

Trust Negotiation Systems <i>Anna Cinzia Squicciarini</i>	90
--------------------------------------------------------------------	----

Advanced Query Processing and Optimization

Continuous Query Processing in Spatio-Temporal Databases <i>Mohamed F. Mokbel</i>	100
Load Distribution for Distributed Stream Processing <i>Ying Xing</i>	112
XML Query Processing and Optimization <i>Ning Zhang</i>	121

Data Models and Database Design

An Access Structure for Similarity Search in Metric Spaces <i>Vlastislav Dohnal</i>	133
Incremental Read-Aheads <i>A. Soydan Bilgin</i>	144
RAM: A Multidimensional Array DBMS <i>Alex R. van Ballegooij</i>	154

Information Retrieval and Interoperability

Handling Inconsistencies in Data Warehouses <i>Mónica Caniupán</i>	166
Data Sharing and Querying for Peer-to-Peer Data Management Systems <i>Anastasios Kementsietsidis</i>	177
Relevance Feedback in XML Retrieval <i>Hanglin Pan</i>	187

Database Technologies for Handling XML-Information on the Web (DataX)

Invited Paper

XML Challenges for the Database Community: Past, Present, and Future <i>Daniele Braga, Alessandro Campi, Stefano Ceri</i>	197
------------------------------------------------------------------------------------------------------------------------------------	-----

Indexing XML Documents

L-Tree: A Dynamic Labeling Structure for Ordered XML Data <i>Yi Chen, George Mihaila, Rajesh Bordawekar, Sriram Padmanabhan</i>	209
Implementation of XPath Axes in the Multi-dimensional Approach to Indexing XML Data <i>Michal Krátký, Jaroslav Pokorný, Václav Snášel</i>	219

Dynamic Range Labeling for XML Trees <i>Takeharu Eda, Yasushi Sakurai, Toshiyuki Amagasa, Masatoshi Yoshikawa, Shunsuke Uemura, Takashi Honishi</i>	230
----------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

<i>FliX: A Flexible Framework for Indexing Complex XML Document Collections</i> <i>Ralf Schenkel</i>	240
-------------------------------------------------------------------------------------------------------------------	-----

Querying XML Documents and Schema Evolution

A Statistical Approach for XML Query Size Estimation <i>Mong Li Lee, Hanyu Li, Wynne Hsu, Beng Chin Ooi</i>	250
--------------------------------------------------------------------------------------------------------------------------	-----

Summarizing XML Data by Means of Association Rules <i>Elena Baralis, Paolo Garza, Elisa Quintarelli, Letizia Tanca</i>	260
-------------------------------------------------------------------------------------------------------------------------------------	-----

Prune XML Before You Search It: XML Transformations for Query Optimization <i>Stéphane Bressan, Zoé Lacroix, Ying Guang Li, Anna Maddalena</i>	270
----------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

Keeping Pace with Evolving XML-Based Specifications <i>Marvin Tan, Angela Goh</i>	280
------------------------------------------------------------------------------------------------	-----

XML Application

Knowledge Management Framework for the Collaborative Distribution of Information <i>Jérôme Godard, Frédéric Andrès, William Grosky, Kinji Ono</i>	289
-------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

XML-Based Revocation and Delegation in a Distributed Environment <i>Konstantina Stoupa, Athena Vakali, Fang Li, Ioannis Tsoukalas</i>	299
----------------------------------------------------------------------------------------------------------------------------------------------------	-----

Using a Transcode and Prefetch Method for Playing XML Contents Containing Multiple Multimedia Data on Mobile Terminals <i>Maria Hong, Dong-Yeop Ryu, Jae-Chul Sir, Eun-Young Kim, Young-Hwan Lim</i>	309
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

Panel

What's Next in XML and Databases? <i>Minos Garofalakis, Ioana Manolescu, Marco Mesiti, George Mihalía, Ralf Schenkel, Bhavani Thuraisingham, Vasilis Vassalos</i>	318
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

Pervasive Information Management (PIM)

Modeling Contextual Information Using Active Data Structures <i>Kevin Goslar, Alexander Schill</i>	325
-----------------------------------------------------------------------------------------------------------------	-----

Context- and Situation-Awareness in Information Logistics <i>Ulrich Meissen, Stefan Pfennigschmidt, Agnès Voisard, Tjark Wahnfried</i>	335
---------------------------------------------------------------------------------------------------------------------------------------------------	-----

An Indexing Scheme for Update Notification in Large Mobile Information Systems <i>Hagen Höpfner, Stephan Schosser, Kai-Uwe Sattler</i>	345
A Mobile Agents Based Architecture for the Distributed Processing of Continuous Location Queries in a Wireless Environment: Performance Evaluation <i>Sergio Ilarri, Eduardo Mena, Arantza Illarramendi</i>	355
Improving Usability of Location-Based Services with User-Centric Data Querying <i>Artem Katasonov</i>	365
MITOS: A Smart Spaces System for Pervasive Computing <i>George Alyfantis, Stathes Hadjiefthymiades, Lazaros Merakos</i>	375
Service Allocation in Selfish Mobile Ad Hoc Networks Using Vickrey Auction <i>Jinshan Liu, Valérie Issarny</i>	385
Engineering Incentive Schemes for Ad Hoc Networks <i>Philipp Obreiter, Birgitta König-Ries, Georgios Papadopoulos</i>	395
The Case for Mobile OLAP <i>Andreas S. Maniatis</i>	405

Peer-to-Peer Computing and Databases (P2P&DB)

Routing

On Constructing Small Worlds in Unstructured Peer-to-Peer Systems <i>Yannis Petrakis, Evaggelia Pitoura</i>	415
SPROUT: P2P Routing with Social Networks <i>Sergio Marti, Prasanna Ganesan, Hector Garcia-Molina</i>	425
A Simulation Framework for Schema-Based Query Routing in P2P-Networks <i>Wolf Siberski, Uwe Thaden</i>	436

Query Evaluation

A Distributed Algorithm for Robust Data Sharing and Updates in P2P Database Networks <i>Enrico Franconi, Gabriel Kuper, Andrei Lopatenko, Ilya Zaihrayeu</i>	446
XPeer: A Self-Organizing XML P2P Database System <i>Carlo Sartiani, Paolo Manghi, Giorgio Ghelli, Giovanni Conforti</i>	456

Bit Zipper Rendezvous Optimal Data Placement for General P2P Queries <i>Wesley W. Terpstra, Stefan Behnel, Ludger Fiege, Jussi Kangasharju, Alejandro Buchmann</i>	466
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

Query Answering in Peer-to-Peer Data Exchange Systems <i>Leopoldo Bertossi, Loreto Bravo</i>	476
-------------------------------------------------------------------------------------------------------	-----

Semantic Query Routing and Processing in P2P Database Systems: The ICS-FORTH SQPeer Middleware <i>George Kokkinidis, Vassilis Christophides</i>	486
-------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

Query Processing in Super-Peer Networks with Languages Based on Information Retrieval: The P2P-DIET Approach <i>Stratos Idreos, Christos Tryfonopoulos, Manolis Koubarakis, Yannis Drougas</i>	496
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----

Data Structures

A Query-Adaptive Partial Distributed Hash Table for Peer-to-Peer Systems <i>Fabius Klemm, Anwitaman Datta, Karl Aberer</i>	506
-------------------------------------------------------------------------------------------------------------------------------------	-----

P2PR-Tree: An R-Tree-Based Spatial Index for Peer-to-Peer Environments <i>Anirban Mondal, Yi Lifu, Masaru Kitsuregawa</i>	516
------------------------------------------------------------------------------------------------------------------------------------	-----

Clustering Information Over the Web (ClustWeb)

Invited Paper

Web Data Protection: Principles and Research Issues <i>Elena Ferrari</i>	526
-----------------------------------------------------------------------------------	-----

Web Content Clustering

Clustering Structured Web Sources: A Schema-Based, Model-Differentiation Approach <i>Bin He, Tao Tao, Kevin Chen-Chuan Chang</i>	536
----------------------------------------------------------------------------------------------------------------------------------------------	-----

Clustering XML Documents Using Structural Summaries <i>Theodore Dalamagas, Tao Cheng, Klaas-Jan Winkel, Timos Sellis</i>	547
-----------------------------------------------------------------------------------------------------------------------------------	-----

A Scalable Randomized Method to Compute Link-Based Similarity Rank on the Web Graph <i>Dániel Fogaras, Balázs Rácz</i>	557
------------------------------------------------------------------------------------------------------------------------------------	-----

Web Logs and Selective Clustering

A Framework for Cluster Management <i>Barbara Catania, Anna Maddalena</i>	568
------------------------------------------------------------------------------------	-----

DiXeminator: A Profile-Based Selective Dissemination System for XML Documents
 Elisa Bertino, Giovanna Guerrini, Marco Mesiti 578

Query Recommendation Using Query Logs in Search Engines
 Ricardo Baeza-Yates, Carlos Hurtado, Marcelo Mendoza 588

Panel

An Overview of Web Data Clustering Practices
 Athena Vakali, Jaroslav Pokorný, Theodore Dalamagas 597

Author Index 607