
Advances in Database Technology — EDBT 2016

19th International Conference
on Extending Database Technology
Bordeaux, France, March 15–18, 2016
Proceedings

Editors

Evaggelia Pitoura
Sofian Maabout
Georgia Koutrika
Amelie Marian
Letizia Tanca
Ioana Manolescu
Kostas Stefanidis



Advances in Database Technology — EDBT 2016
Proceedings of the 19th International Conference
on Extending Database Technology
Bordeaux, France, March 15–18, 2016

Series ISSN: 2367-2005

Editors

Evaggelia Pitoura, University of Ioannina, Greece
Sofian Maabout, University of Bordeaux, France
Georgia Koutrika, HP Labs, USA
Amelie Marian, Rutgers University, USA
Letizia Tanca, Politecnico di Milano, Italy
Ioana Manolescu, INRIA, France
Kostas Stefanidis, ICS-FORTH, Greece



OpenProceedings.org
University of Konstanz
University Library
78457 Konstanz, Germany

COPYRIGHT NOTICE: Copyright © 2016 by the authors of the individual papers.

Distribution of all material contained in this volume is permitted under the terms of the Creative Commons license
CC-by-nc-nd 4.0

OpenProceedings ISBN: 978-3-89318-070-7

DOI of this front matter: 10.5441/002/edbt.2016.01

Foreword

Welcome to the 19th edition of the International Conference of Extending Database Technology (EDBT). Originally biennial, the EDBT conference has been held annually and jointly with ICDT (“International Conference on Database Theory”) since 2009. This year, EDBT is taking place in Bordeaux, France, on March 15–18, 2016, continuing its long tradition as a top venue for presenting and discussing recent advancements in data management.

This year we received 137 submissions to the research track, 18 submissions to the vision track, 24 submissions to the industrial/application track, 28 submissions to the demo track and 11 tutorial proposals. The high quality of these submissions made the job of selecting the best of them rather challenging. The various program committees after thorough reviewing and careful consideration selected 38 research papers, 5 vision papers, 9 industrial/application papers, 16 demos and 3 tutorials. The proceedings include these contributions. A new addition this year is the poster track for presenting novel ongoing work. There were 64 submissions from which the poster program committee selected 31 contributions included in this proceedings.

The proceedings also include an overview of the keynote talk by Elisa Bertino (Purdue), an overview of the keynote talk by Gustavo Alonso (ETHZ) and a laudation concerning the EDBT 2016 Test of Time Award that was given to the paper

“Bridging Physical and Virtual Worlds: Complex Event Processing for RFID Data Streams” by Fusheng Wang, Shaorong Liu, Peiya Liu, Yijian Bai, published in the EDBT 2006 proceedings.

The EDBT 2016 program is the result of the joint effort of many people that I would like to take this opportunity to thank. Ioana Manolescu (Vision Track Chair), Georgia Koutrika (Industrial/Application Track Chair), Letizia Tanca (Demo Track Chair) and Amelie Marian (Tutorial Chair), all did an excellent job, as Themis Palpanas with the workshops (whose proceedings appear in a companion volume). Thanks also to the members of the program committees of the various tracks that worked very hard to review each submission in detail and engaged in many discussions to create the best possible program.

Special mention should be made to the Test of Time Award committee members: Sihem Amer-Yahia, Yannis Ioannidis and Christian S. Jensen. The general chair, Sofian Maabout and the local organizers worked hard with all arrangements necessary for securing a successful event. Special thanks to Kostas Stefanidis, the proceedings chair, and Patrick Mary, the website chair, for their invaluable contribution to this event. Christine Collet and Norman Paton were instrumental in advising and coordinating with the EDBT Executive Board.

And lastly and most importantly, thanks to all the authors that submitted their work to EDBT 2016. Their contributions were what made this a strong program. I hope that you find the EDBT 2016 conference informative, enjoyable and thought-provoking!

Evaggelia Pitoura
EDBT 2016 Program Chair

Program Committee Members

Research Program Committee

Bernd Amann (U Pierre et Marie Curie)
Walid Aref (Purdue U)
Sourav S Bhowmick (Nanyang TU)
Michael Böhlen (U of Zurich)
Klemens Böhm (KIT)
Francesco Bonchi (Yahoo! Labs)
Angela Bonifati (Lille 1 U)
Philippe Bonnet (ITU)
Luc Bouganim (INRIA)
Nieves Brisaboa (U de La Coruna)
Reynold Cheng (U of Hong Kong)
Beng Chin Ooi (National U of Singapore)
Vassilis Christophides (INRIA Paris)
Panos K Chrysanthis (U of Pittsburgh)
Paolo Ciaccia (U of Bologna)
Philippe Cudre-Mauroux (U of Fribourg)
Bin Cui (Peking U)
Alfredo Cuzzocrea (U of Trieste)
Khuzaima Daudjee (U of Waterloo)
Antonios Deligiannakis (TU of Crete)
Elena Ferrari (U of Insubria)
Peter Fischer (U Freiburg)
Helena Galhardas (U of Lisbon)
Johann Gamper (Free U Bolzano)
Minos Garofalakis (TU of Crete)
Floris Geerts (U of Antwerp)
Jiawei Han (UI Urbana Champaign)
Takahiro Hara (Osaka U)
Thomas Heinis (Imperial College)
Arantza Illarramendi (U del Paes Vasco)
George Kollios (Boston U)
Georgia Koloniari (U of Macedonia)
Yiannis Kotidis (Athens U of Bus. & Econ.)
Nick Koudas (U of Toronto)
Georg Lausen (U Freiburg)
Wang-Chien Lee (Penn State U)
Wolfgang Lehner (TU Dresden)
Hong-Va Leong (Hong Kong Polytechnic U)
Roy Levin (IBM Research)
Feifei Li (U of Utah)
Xuemin Lin (U of New South Wales)
Eric Lo (Honk Kong Polytechnic)
Norman May (SAP)
Sebastian Michel (TU Kaiserslautern)
Kjetil Norvag (Norwegian U of Sc. & Tech.)
Ippokratis Pandis (Cloudera)
Paolo Papotti (QCRI)
Marta Patino (Politecnico de Madrid)
Torben B Pedersen (U of Aalborg)
Peter Pietzuch (Imperial College)
Maya Ramanath (IIT Delhi)
Matthias Renz (LMU)
Rodolfo Resende (U Federal de Minas Gerais)
Tore Risch (Uppsala U)
Pierangela Samarati (U Studi Milano)
Mohamed Sarwat (Arizona State U)
Kai-Uwe Sattler (TU Ilmenau)
Marc Scholl (U of Konstanz)
Heiko Schuldt (U of Basel)
Assaf Schuster (Technion)
Thomas Seidl (RWTH Aachen)
Jianwen Su (UC Santa Barbara)
Peter Triantafillou (U of Glasgow)
Yannis Velegarakis (U of Trento)
Stratis Viglas (U of Edinburgh)
Jef Wijsen (U of Mons – UMONS)
Yoshitaka Yamamoto (U of Yamanashi)
Carlo Zaniolo (UCLA)
Demetrios Zeinalipour-Yazti (U of Cyprus)
Wenjie Zhang (U of New South Wales)

Vision Track Committee

Nicolas Ancaux (INRIA Paris-Rocquencourt)
Iovka Boneva (U Lille 1)
Yanlei Diao (Ecole Polytechnique)
Stratos Idreos (Harvard U)
Yannis Ioannidis (U of Athens)
Christian Jensen (Aalto U)
Alekh Jindal (Microsoft)
Zoi Kaoudi (QCRI)
Giansalvatore Mecca (U della Basilicata)
Leonid Libkin (U of Edinburgh)
Neoklis Polyzotis (Google)
Nicoleta Preda (U de Versailles)
Eric Simon (SAP)
Alessandro Solimando (INRIA)
Fabian Suchanek (Télécom ParisTech)

Industrial Program Committee

Andrey Balmin (Platfora)
Fei Chen (HP Labs)
Vuk Ercegovic (Google)
Mohamed Eltabakh (Worcester PI)
Irina Fundulaki (ICS-FORTH)
Oktie Hassanzadeh (IBM Watson)
Anastasios Kementsietsidis (Google)
Lipyew Lim (U of Hawaii)
Konstantinos Morfonios (Oracle)
Lucian Popa (IBM Almaden Research)
Lin Qiao (LinkedIn)
Mohamed Sharaf (U of Queensland)
Julia Stoyanovich (Drexel U)
Nesime Tatbul (Intel Labs and MIT)
Panayiotis Tsaparas (U of Ioannina)
Steven (Euijong) Whang (Google)
Kevin Wilkinson (HP)

Poster Track Committee

Alberto Abelló (Politécnica de Catalunya)
Nikolaus Augsten (U of Salzburg)
Christos Doulkeridis (U of Piraeus)
Ioana Giurgiu (IBM Research (Zurich))
Aris Gkoulalas-Divanis (IBM Research)
Sven Groppe (U of Lubeck)
Katja Hose (Aalborg U)
Verena Kantere (U of Geneva)
Viktor Leis (Technische Ut Munchen)
Paolo Missier (Newcastle U)
Eirini Ntoutsis (LMU)
Senjuti Basu Roy (U of Washington Tacoma)
George Pallis (U of Cyprus)
Shaoyu Song (Tsinghua U)

External Reviewers

Daichi Amagata (Osaka U)
Mohammad Amiri (UC Santa Barbara)
Khaled Ammar (U of Waterloo)
Christos Anagnostopoulos (U of Glasgow)

Carlos Andrade (U of Hawaii at Manoa)
Ilaria Bartolini (U di Bologna)
Dritan Bleco (AUEB)
Carlos Bobed (U of Zaragoza)
Douglas Burdick (IBM Research Almaden)
Siarhei Bykau (Purdue U)
Lijun Chang (UNSW)
Georgios Chatzimilioudis (U of Cyprus)
Sean Chester (NTNU)
Pietro Colombo (U of Insubria)
Camelia Constantin (U P&M Curie)
Maria Daltayanni (U of San Francisco)
Vasilis Efthymiou (U of Crete)
Ioanna Filippidou (AUEB)
George Fletcher (Eindhoven UT)
Sara Foresti (U degli Studi di Milano)
Daniele Foroni (U of Trento)
Shi Gao (UCLA)
Xiaoyu (Steve) Ge (U of Pittsburgh)
Kostas Georgoulas (AUEB)
Orestis Gkorgkas (NTNU)
Alfredo Goni (Basque Country U)
Zengfeng Huang (UNSW)
Meng Jiang (UIUC)
Julius Koepke (U of Klagenfurt)
Mustafa Korkmaz (U of Waterloo)
Zeynep Korkmaz (U of Waterloo)
Christos Laoudias (U of Cyprus)
Jialu Liu (UIUC)
Giovanni Livraga (U Milano)
Xiuli Ma (Peking U)
Massimo Mazzeo (UCLA)
Evica Milchevski (TU Kaiserslautern)
Davide Mottin (U of Trento)
Hubert Naacke (UPMC-LIP6)
Nathan Rico Ong (U of Pittsburgh)
Kiril Panev (TU Kaiserslautern)
Marco Patella (U di Bologna)
Fabio Petroni (Sapienza U of Rome)
Yoann Pitarch (U Paul Sabatier)
Donatello Santoro (U della Basilicata)
Klaus Schmid (LMU)
Konstantinos Semertzidis (U of Ioannina)
Anatoli Shein (U of Pittsburgh)
Masumi Shirakawa (Osaka U)
Vasilis Spyropoulos (AUEB)
Yan Tang (UC Santa Barbara)
Io Taxidou (U of Freiburg)
Cory Thoma (U of Pittsburgh)
Sabrina De Capitani di Vimercati (U Milano)
Xiaoyang Wang (UNSW)
Doris Xin (UIUC)
Mohan Yang (UCLA)
Man Lung Yiu (Hong Kong Polytechnic U)
Quan Yuan (UIUC)
Roberto Yus (U of Zaragoza)
Chao Zhang (UIUC)
Andreas Zuefle (LMU)

Test-of-Time Award

In 2014, EDBT began awarding the EDBT Test-of-Time (ToT) Award, with the goal of recognizing one paper, or a small number of papers, presented at EDBT earlier and that have best met the “test of time”, i.e., that has had the most impact in terms of research, methodology, conceptual contribution, or transfer to practice over the past decade(s). The EDBT ToT Award for 2016 will be presented during the EDBT/ICDT 2016 Joint Conference, March 15–18, 2016, in Bordeaux (France). The EDBT 2016 Test-of-Time Award committee was formed by Sihem Amer-Yahia (CNRS, Laboratoire d’Informatique de Grenoble, France), Yannis Ioannidis (University of Athens, Greece), Christian S. Jensen (Aalborg University, Denmark), and all PC chairs of former EDBT conferences including EDBT 2006.

The committee was asked to select a paper or a small number of papers from the EDBT 2006 (Munich) proceedings. After careful consideration, the committee and the EDBT Executive Board have decided to select the following paper as the EDBT ToT Award winner for 2016:

**Bridging Physical and Virtual Worlds:
Complex Event Processing for RFID Data Streams**

by Fusheng Wang, Shaorong Liu, Peiya Liu, Yijian Bai

published in the EDBT 2006 proceedings, 588–607

The paper proposes an event-oriented approach to the processing of RFID data which makes it possible to automate the translation of RFID based application semantics through complex event detection. In particular, it demonstrates the ability to process complex events by capturing temporal constraints in an algebra. The resulting declarative event-based approach is shown to simplify RFID data processing and is shown to be scalable. The paper pioneers declarative event-based RFID processing. The simplicity and expressiveness of the proposed framework are admirable. For example, the framework makes it possible to express object tracking on historical data as well as to formulate real-time monitoring.

The committee and the EDBT Executive Board find that this paper stands out in terms of relevance, impact, and influence in databases. It has had substantial impact. In particular, it has impacted real systems, and the engine it proposes has been integrated into Siemens RFID Middleware. It is also the most cited EDBT 2006 paper, has spurred a significant amount of follow-up work, and remains relevant today.

Table of Contents

Foreword	i
Program Committee Members	ii
Test-of-Time Award	iv
Table of Contents	v
Invited Keynotes	1
Data Security and Privacy in the IoT <i>Elisa Bertino</i>	1
Data Processing in Modern Hardware <i>Gustavo Alonso</i>	4
Research Papers	5
Finding Users of Interest in Micro-blogging Systems <i>Camelia Constantin, Ryadh Dahimene, Quentin Grossetti, Cedric Du Mouza</i>	5
Slowing the Firehose: Multi-Dimensional Diversity on Social Post Streams <i>Shiwen Cheng, Marek Chrobak, Vagelis Hristidis</i>	17
Social, Structured and Semantic Search <i>Raphaël Bonaque, Bogdan Cautis, François Goasdoué, Ioana Manolescu</i>	29
Indexing Query Graphs to Speedup Graph Query Processing <i>Jing Wang, Nikos Ntarmos, Peter Triantafillou</i>	41
GSCALER: Synthetically Scaling A Given Graph <i>J.W. Zhang, Y.C. Tay</i>	53
Storing and Analyzing Historical Graph Data at Scale <i>Udayan Khurana, Amol Deshpande</i>	65
Providing Serializability for Pregel-like Graph Processing Systems <i>Minyang Han, Khuzaima Daudjee</i>	77
DBExplorer: Exploratory Search in Databases <i>Manish Singh, Michael Cafarella, Hosagrahar Visvesvar Jagadish</i>	89
Refinement Driven Processing of Aggregation Constrained Queries <i>Manasi Vartak, Venkatesh Raghavan, Elke Rundensteiner, Samuel Madden</i>	101
Reverse Engineering Top-k Database Queries with PALEO <i>Kiril Panev, Sebastian Michel</i>	113
CrowdSky: Skyline Computation with Crowdsourcing <i>Jongwuk Lee, Dongwon Lee, Sang-Wook Kim</i>	125
Cohesive Keyword Search on Tree Data <i>Aggeliki Dimitriou, Ananya Dass, Dimitri Theodoratos, Yannis Vassiliou</i>	137
Generic Keyword Search over XML Data <i>Manoj Agarwal, Krithi Ramamritham, Prashant Agarwal</i>	149
Answering Keyword Queries involving Aggregates and GROUPBY on Relational Databases <i>Zhong Zeng, Mong Li Lee, Tok Wang Ling</i>	161

Finding All Maximal Cliques in Very Large Social Networks <i>Alessio Conte, Roberto De Virgilio, Antonio Maccioni, Maurizio Patrignani, Riccardo Torlone</i>	173
RPM: Representative Pattern Mining for Efficient Time Series Classification <i>Xing Wang, Jessica Lin, Pavel Senin, Tim Oates, Sunil Gandhi, Arnold Boedihardjo, Crystal Chen, Susan Frankenstein</i>	185
Interactive Temporal Association Analytics <i>Xiao Qin, Ramoza Ahsan, Xika Lin, Elke Rundensteiner, Matthew Ward</i>	197
Efficient Record Linkage Using a Compact Hamming Space <i>Dimitrios Karapiperis, Dinusha Vatsalan, Vassilios Verykios, Peter Christen</i>	209
Scaling Entity Resolution to Large, Heterogeneous Data with Enhanced Meta-blocking <i>George Papadakis, George Papastefanatos, Themis Palpanas, Manolis Koubarakis</i>	221
Practical Query Answering in Data Exchange Under Inconsistency-Tolerant Semantics <i>Balder ten Cate, Richard Halpert, Phokion Kolaitis</i>	233
Querying RDF Data Using A Multigraph-based Approach <i>Vijay Ingalalli, Dino Ienco, Pascal Poncelet, Serena Villata</i>	245
Optimization of Complex SPARQL Analytical Queries <i>Padmashree Ravindra, HyeongSik Kim, Kemafor Anyanwu</i>	257
RDF-TX: A Fast, User-Friendly System for Querying the History of RDF Knowledge Bases <i>Shi Gao, Jiaqi Gu, Carlo Zaniolo</i>	269
Efficient Computation of Containment and Complementarity in RDF Data Cubes <i>Marios Meimaris, George Papastefanatos, Panos Vassiliadis, Ioannis Anagnostopoulos</i>	281
Semi-automatic support for evolving functional dependencies <i>Mirjana Mazuran, Elisa Quintarelli, Letizia Tanca, Stefania Ugolini</i>	293
Holistic Data Profiling: Simultaneous Discovery of Various Metadata <i>Jens Ehrlich, Mandy Roick, Lukas Schulze, Jakob Zwiener, Thorsten Papenbrock, Felix Naumann</i>	305
Monitoring MaxRS in Spatial Data Streams <i>Daichi Amagata, Takahiro Hara</i>	317
Similarity Search on Spatio-Textual Point Sets <i>Christodoulos Efstathiades, Alexandros Belesiotis, Dimitrios Skoutas, Dieter Pfoser</i>	329
Nearest Window Cluster Queries <i>Chen-Che Huang, Jiun-Long Huang, Tsung-Ching Liang, Jun-Zhe Wang, Wen-Yuah Shih, Wang-Chien Lee</i>	341
Adaptive query parallelization in multi-core column stores <i>Mrunal Gawade, Martin Kersten</i>	353
PARAGON: Parallel Architecture-Aware Graph Partition Refinement Algorithm <i>Angen Zheng, Alexandros Labrinidis, Patrick Pisciuneri, Panos Chrysanthis, Peyman Givi</i>	365
Query Workload-based RDF Graph Fragmentation and Allocation <i>Peng Peng, Lei Zou, Lei Chen, Dongyan Zhao</i>	377
Efficient Query Processing using the Earth's Mover Distance in Video Databases <i>Merih Seran Uysal, Christian Becks, Daniel Sabinasz, Jochen Schmuecking, Thomas Seidl</i>	389
Probabilistic Threshold Indexing for Uncertain Strings <i>Sudip Biswas, Manish Patil, Sharma Thankachan, Rahul Shah</i>	401

Context-Aware Event Stream Analytics <i>Olga Poppe, Chuan Lei, Elke Rundensteiner, Dan Dougherty</i>	413
Who Cares about Others' Privacy: Personalized Anonymization of Moving Object Trajectories <i>Despina Kopanaki, Vasilis Theodossopoulos, Nikos Pelekis, Ioannis Kopanakis, Yannis Theodoridis</i> . . .	425
Identifying and Describing Streets of Interest <i>Dimitrios Skoutas, Dimitris Sacharidis, Kostas Stamatoukos,</i>	437
Finding Frequently Visited Indoor POIs Using Symbolic Indoor Tracking Data <i>Hua Lu, Chenjuan Guo, Bin Yang, Christian Jensen</i>	449
Visionary Papers	461
Designing Access Methods: The RUM Conjecture <i>Manos Athanassoulis, Michael Kester, Lukas Maas, Radu Stoica, Stratos Idreos, Anastasia Ailamaki, Mark Callaghan</i>	461
Self-Curating Databases <i>Mohammad Sadoghi, Kavitha Srinivas, Oktie Hassanzadeh, Yuan-Chi Chang, Mustafa Canim, Achille Fokoue, Yishai Feldman</i>	467
Data Wrangling for Big Data: Challenges and Opportunities <i>Tim Furge, Georg Gottlob, Leonid Libkin, Giorgio Orsi, Norman Paton</i>	473
Road to Freedom in Big Data Analytics <i>Divy Agrawal, Sanjay Chawla, Ahmed Elmagarmid, Zoi Kaoudi, Mourad Ouzzani, Paolo Papotti, Jorge Quiane, Nan Tang, Mohammed Zaki</i>	479
Data Management for Next Generation Genomic Computing <i>Stefano Ceri, Abdulrahman Kaitoua, Marco Masseroli, Pietro Pinoli, Francesco Venco</i>	485
Industrial and Applications Papers	491
Exploring Text Classification for Messy Data: An Industry Use Case for Domain-Specific Analytics <i>Laura Kassner, Bernhard Mitschang</i>	491
Discovering Correlations in Annotated Databases <i>Xuebin He, Stephen Donohue, Mohamed Eltabakh</i>	503
Query Performance Problem Determination with Knowledge Base in Semantic Web System OptImatch <i>Guilherme Damasio, Piotr Mierzejewski, Jaroslaw Szlichta, Calisto Zuzarte</i>	515
Scalable Public Transportation Queries on the Database <i>Alexandros Efentakis</i>	527
Characterizing Home Device Usage From Wireless Traffic Time Series <i>Katsiaryna Mirylenka, Vassilis Christophides, Themis Palpanas, Ioannis Pefkianakis, Martin May</i> . . .	539
Parallel Duplicate Detection in Adverse Drug Reaction Databases with Spark <i>Chen Wang, Sarvnaz Karimi</i>	551
e#: Sharper Expertise Detection from Microblogs <i>Thibault Sellam, Martin Hentschel, Vasilis Kandylas, Omar Alonso</i>	563
DECT: Distributed Evolving Context Tree for Mining Web Behavior Evolution <i>Xiaokui Shu, Nikolay Laptev, Danfeng Yao</i>	573
Strudel: A Framework for Transaction Performance Analyses on SQL/NoSQL Systems <i>Junichi Tatemura, Oliver Po, Zheng Li, Hakan Hacigumus</i>	580

Demonstrations	592
GROM: a General Rewriter of Semantic Mappings <i>Giansalvatore Mecca, Guillem Rull, Donatello Santoro, Ernest Teniente</i>	592
PowerQ: An Interactive Keyword Search Engine for Aggregate Queries on Relational Databases <i>Zhong Zeng, Mong Li Lee, Tok Wang Ling</i>	596
Visualization Through Inductive Aggregation <i>Parke Godfrey, Jarek Gryz, Piotr Lasek, Nasim Razavi</i>	600
Contextual Event Search: Finding Contextual Messages in Dynamic microblog Data Stream in Real Time <i>Manoj Agarwal, Divyam Bansal, Mridul Garg, Krithi Ramamritham</i>	604
Answering Controlled Natural Language Questions on RDF Knowledge Bases <i>Giuseppe Mazzeo, Carlo Zaniolo</i>	608
tPredictor: A Micro-blog Based System for Teenagers' Stress Prediction <i>Jing Huang, Qi Li, Zhuonan Feng, Yiping Li, Ling Feng</i>	612
OSNI: Searching for Needles in a Haystack of Social Network Data <i>Shiwen Cheng, James Fang, Vagelis Hristidis, Harsha Madhyastha, Niluthpol Chowdhury Mithun, Dorian Perkins, Amit Roy-Chowdhury, Moloud Shahbazi, Vassilis Tsotras</i>	616
PROX: Approximated Summarization of Data Provenance <i>Eleanor Ainy, Pierre Bourhis, Susan Davidson, Daniel Deutch, Tova Milo</i>	620
PAW: A Platform for Analytics Workflows <i>Maxim Filatov, Verena Kantere</i>	624
streamLoader: An Event-Driven ETL System for the On-line Processing of Heterogeneous Sensor Data <i>Marco Mesiti, Luca Ferrari, Stefano Valtolina, Giacomo Licari, Gianluca Galliani, Minh-San Dao, Koji Zettsu</i>	628
TINTIN: a Tool for INcremental INtegrity checking of Assertions in SQL Server <i>Xavier Oriol, Ernest Teniente, Guillem Rull</i>	632
Efficient regular path query evaluation using path indexes <i>George Fletcher, Jeroen Peters, Alexandra Poulouvassilis</i>	636
Galaxy: A Platform for Explorative Analysis of Open Data Sources <i>Seyed-Mehdi-Reza Beheshti, Boualem Benatallah, Hamid Reza Motahari Nezhad,</i>	640
OAPT: A Tool for Ontology Analysis and Partitioning <i>Alsayed Algergawy, Samira Babalou, Friederike Klan, Birgitta König-Ries</i>	644
ShapeExplorer: Querying and Exploring Shapes using Visual Knowledge <i>Tong Ge, Yafang Wang, Gerard de Melo, Zengguang Hao, Andrei Sharf, Baoquan Chen</i>	648
Distributed Secure Search in the Personal Cloud <i>Thu Le, Nicolas Anceaix, Sebastien Guilloton, Saliha Lallali, Philippe Pucheral, Iulian Sandu Popa, Chao Chen</i>	652
Poster Papers	656
Type-aware Web-search <i>Michael Gubanov, Anna Pyayt</i>	656
Indexing and Querying A Large Database of Typed Intervals <i>Jianqiu Xu, Hua Lu, Bin Yao</i>	658
Quantifying Likelihood of Change through Update Propagation across Top-k Rankings <i>Evica Milchevski, Sebastian Michel</i>	660

Optimizing B+-Tree for PCM-Based Hybrid Memory <i>Lu Li, Peiquan Jin, Chengcheng Yang, Zhanglin Wu, Lihua Yue</i>	662
A Data Mining Approach to Choosing Categorical Attributes for Ranked Lists <i>Koninika Pal, Sebastian Michel</i>	664
Efficient Implementation of Joins over Cassandra DBs <i>Haridimos Kondylakis, Antonis Fountouris, Dimitris Plexousakis</i>	666
Double Chain-Star: an RDF indexing scheme for fast processing of SPARQL joins <i>Marios Meimaris, George Papastefanatos</i>	668
Minoan ER: Progressive Entity Resolution in the Web of Data <i>Vasilis Efthymiou, Kostas Stefanidis, Vassilis Christophides</i>	670
Proposal of a Database Type and Aggregation Function for Accelerating Medical Genomics Study on RDBMS <i>Yoshifumi Ujibashi, Motoyuki Kawaba, Lilian Harada</i>	672
The Best Bang for Your Bu(ck)g <i>Benjamin Dietrich, Tobias Müller, Torsten Grust</i>	674
A Way to Automatically Enrich Biomedical Ontologies <i>Juan Antonio Lossio-Ventura, Mathieu Roche, Clement Jonquet, Maguelonne Teisseire</i>	676
A Distributed Mining Framework for Influence in Evolving Entities <i>Tian Guo, Karl Aberer</i>	678
Sweet KIWI: Statistics-Driven OLAP Acceleration using Query Column Sets <i>Sung-Soo Kim, Taewhi Lee, Moonyoung Chung, Jongho Won</i>	680
On-Line Mobility Pattern Discovering using Trajectory Data <i>Ticiansia Coelho da Silva, Karine Zeitouni, José Fernandes de Macêdo, Marco Casanova</i>	682
Summarizing Linked Data RDF Graphs Using Approximate Graph Pattern Mining <i>Mussab Zneika, Claudio Lucchese, Dan Vodislav, Dimitris Kotzinos</i>	684
Understanding Customer Attrition at an Individual Level: a New Model in Grocery Retail Context <i>Clément Gautrais, Peggy Cellier, Thomas Guyet, René Quiniou, Alexandre Termier</i>	686
Towards an Efficient Ranking of Interval-Based Patterns <i>Marwan Hassani, Yifeng Lu, Thomas Seidl</i>	688
SOFYA: Semantic on-the-fly Relation Alignment <i>Maria Koutraki, Nicoleta Preda, Dan Vodislav</i>	690
Model Kit for Lightweight Data Compression Algorithms <i>Juliana Hildebrandt, Dirk Habich, Patrick Damme, Wolfgang Lehner</i>	692
Revisiting DBMS Space Management for Native Flash <i>Sergey Hardock, Iliia Petrov, Robert Gottstein, Alejandro Buchmann</i>	694
A Two Phase Deep Learning Model for Identifying Discrimination from Tweets <i>Shuhan Yuan, Xintao Wu, Yang Xiang</i>	696
Top-k Dominating Queries, in Parallel, in Memory <i>Sean Chester, Orestis Gkorgkas, Kjetil Nørkvåg</i>	698
Snapshot Isolation for Neo4j <i>Marta Patino, Ricardo Jimenez-PEris, Diego Burgos-Sancho, Ivan Brondino, Valerio Vianello, Rohit Dhamane</i>	700

Maximum Coverage Representative Skyline <i>Malene S�holm, Sean Chester, Ira Assent</i>	702
An On-Line Approximation Algorithm for Mining Frequent Closed Itemsets Based on Incremental Inter- section <i>Koji Iwanuma, Yoshitaka Yamamoto, Shoshi Fukuda</i>	704
Extending Database Accelerators for Data Transformations and Predictive Analytics <i>Felix Beier, Knut Stolze, Daniel Martin</i>	706
Privacy Protection through Query Rewriting in Smart Environments <i>Hannes Grunert, Andreas Heuer</i>	708
DatShA :A Data Sharing Algebra for access control plans <i>Luc Bouganim, Athanasia Katsouraki, Benjamin Nguyen</i>	710
Cluster-based Contextual Recommendations <i>Kostas Stefanidis, Eirini Ntoutsis</i>	712
Empirical evaluation of guarded structural indexing <i>Erik Agterdenbos, George Fletcher, Chee-Yong Chan, Stijn Vansummeren</i>	714
Context-Dependent Quality-Aware Source Selection for Live Queries on Linked Data <i>Barbara Catania, Giovanna Guerrini, Beyza Yaman</i>	716
Tutorials	718
Data Responsibly: Fairness, Neutrality and Transparency in Data Analysis <i>Julia Stoyanovich, Serge Abiteboul, Jerome Miklau</i>	718
Core Decomposition in Graphs: Concepts, Algorithms and Applications <i>Fragkiskos D. Malliaros, Apostolos N. Papadopoulos, Michalis Vazirgiannis</i>	720
Distance-based Multimedia Indexing <i>Christian Becks, Merih Seran Uysal, Thomas Seidl</i>	722