

19th International Conference on Database Theory

ICDT'16, March 15–18, 2016, Bordeaux, France

Edited by

Wim Martens

Thomas Zeume



Editors

Wim Martens	Thomas Zeume
University of Bayreuth	TU Dortmund University
Bayreuth, Germany	Dortmund, Germany
wim.martens@uni-bayreuth.de	thomas.zeume@tu-dortmund.edu

ACM Classification 1998

H.2 Database Management, H.2.1 Normal forms, H.2.2 Schema and subschema, H.2.3 Query languages, H.2.4 Query processing, H.2.4 Relational databases, H.2.4 Distributed databases, H.2.5 Heterogeneous Databases, H.3.5 Online Information Services, H.1 Miscellaneous – Privacy, H.4.1 Office Automation: Workflow management, B.4.4 Performance Analysis and Design Aids: Formal models, Verification, F.1.3 Complexity measures and classes, F.4.1 Computational Logic, Model Theory, G.2.2 Graph Theory|Hypergraphs

ISBN 978-3-95977-002-6

Published online and open access by

Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at <http://www.dagstuhl.de/dagpub/978-3-95977-002-6>.

Publication date

March, 2016

Bibliographic information published by the Deutsche Nationalbibliothek

The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at <http://dnb.d-nb.de>.

License

This work is licensed under a Creative Commons Attribution 3.0 Unported license (CC-BY 3.0): <http://creativecommons.org/licenses/by/3.0/legalcode>.



In brief, this license authorizes each and everybody to share (to copy, distribute and transmit) the work under the following conditions, without impairing or restricting the authors' moral rights:

- Attribution: The work must be attributed to its authors.

The copyright is retained by the corresponding authors.

Digital Object Identifier: 10.4230/LIPICS.ICDT.2016.0

ISBN 978-3-95977-002-6

ISSN 1868-8969

<http://www.dagstuhl.de/lipics>

Contents

Preface	0:ix
Organization	0:xi
External Reviewers	0:xiii
List of Authors	0:xv

ICDT 2016 Test of Time Award

The ICDT 2016 Test of Time Award Announcement	
<i>Foto N. Afrati, Claire David, and Georg Gottlob</i>	1:1–1:2

Invited Talks

Scale Independence: Using Small Data to Answer Queries on Big Data	
<i>Floris Geerts</i>	2:1–2:2
Top-k Indexes Made Small and Sweet	
<i>Yufei Tao</i>	3:1–3:1
New Algorithms for Heavy Hitters in Data Streams	
<i>David P. Woodruff</i>	4:1–4:12

Awards Session

Beyond Well-designed SPARQL	
<i>Mark Kaminski and Egor V. Kostylev</i>	5:1–5:18
A Framework for Estimating Stream Expression Cardinalities	
<i>Anirban Dasgupta, Kevin J. Lang, Lee Rhodes, and Justin Thaler</i>	6:1–6:17
(Regular Paper)	
Declarative Probabilistic Programming with Datalog	
<i>Vince Barany, Balder ten Cate, Benny Kimelfeld, Dan Olteanu, and Zografoula Vagena</i>	7:1–7:19

Distribution and Paralellism

Worst-Case Optimal Algorithms for Parallel Query Processing	
<i>Paraschos Koutris, Paul Beame, and Dan Suciu</i>	8:1–8:18
Parallel-Correctness and Containment for Conjunctive Queries with Union and Negation	
<i>Gaetano Geck, Bas Ketsman, Frank Neven, and Thomas Schwentick</i>	9:1–9:17

19th International Conference on Database Theory (ICDT 2016).

Editors: Wim Martens and Thomas Zeume



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

A Formal Study of Collaborative Access Control in Distributed Datalog <i>Serge Abiteboul, Pierre Bourhis, and Victor Vianu</i>	10:1–10:17
---	------------

Optimization

It’s All a Matter of Degree: Using Degree Information to Optimize Multiway Joins <i>Manas R. Joglekar and Christopher M. Ré</i>	11:1–11:17
Filtering With the Crowd: CrowdScreen Revisited <i>Benoît Groz, Ezra Levin, Isaac Meilijson, and Tova Milo</i>	12:1–12:18
Streaming Partitioning of Sequences and Trees <i>Christian Konrad</i>	13:1–13:18

Evolving Data

Dynamic Graph Queries <i>Pablo Muñoz, Nils Vortmeier, and Thomas Zeume</i>	14:1–14:18
Verification of Evolving Graph-structured Data under Expressive Path Constraints <i>Diego Calvanese, Magdalena Ortiz, and Mantas Šimkus</i>	15:1–15:19
Query Stability in Monotonic Data-Aware Business Processes <i>Ognjen Savković, Elisa Marengo, and Werner Nutt</i>	16:1–16:18

Data Extraction and Analytics

Document Spanners: From Expressive Power to Decision Problems <i>Dominik D. Freydenberger and Mario Holldack</i>	17:1–17:17
Algorithms for Provisioning Queries and Analytics <i>Sepehr Assadi, Sanjeev Khanna, Yang Li, and Val Tannen</i>	18:1–18:18

Schemas and Consistency

Limits of Schema Mappings <i>Phokion G. Kolaitis, Reinhard Pichler, Emanuel Sallinger, and Vadim Savenkov</i> .	19:1–19:17
Reasoning About Integrity Constraints for Tree-structured Data <i>Wojciech Czerwiński, Claire David, Filip Murlak, and Paweł Parys</i>	20:1–20:18
Complexity of Repair Checking and Consistent Query Answering <i>Sebastian Arming, Reinhard Pichler, and Emanuel Sallinger</i>	21:1–21:18

SPARQL and Graphs

On the Complexity of Enumerating the Answers to Well-designed Pattern Trees <i>Markus Kröll, Reinhard Pichler, and Sebastian Skritek</i>	22:1–22:18
---	------------

A Practically Efficient Algorithm for Generating Answers to Keyword Search Over Data Graphs <i>Konstantin Golenberg and Yehoshua Sagiv</i>	23:1–23:17
--	------------