



Special Issue on Database Theory

Pablo Barceló¹ · Marco Calautti²

Published online: 27 August 2020

© Springer Science+Business Media, LLC, part of Springer Nature 2020

This volume is dedicated to invited papers from the 22nd edition of the International Conference on Database Theory (ICDT 2019), held in Lisbon, Portugal, on March 26–29, 2019. The ICDT conference is one of the leading venues in database theory and foundations of data management. This is an exciting area at the core of efficient and effective data management, with further connections to knowledge representation, relational statistical learning, logical aspects of computation, verification, and, of course, practical aspects of data processing.

Based on the results of the reviewing process and oral presentations given at the conference, a group of program committee members from ICDT 2019 selected five articles from the conference to be invited to this special issue. These are among the finest contributions from ICDT 2019. The invited papers were further reviewed according to the journal's rigorous peer-review standards.

The five selected papers deal with different, yet timely problems in the area of foundations of data management. In particular:

- The paper “Characterizing Tractability of Simple Well-designed Pattern Trees with Projection”, by Stefan Mengel and Sebastian Skritek, deals with the challenging problem of characterizing which well-designed pattern trees can be evaluated efficiently. Such well-designed pattern trees lie at the core of modern languages for semantic web and graph databases.
- The paper “Index-Based, High-Dimensional, Cosine Threshold Querying with Optimality Guarantees”, by Yuliang Li Janguo Wang Benjamin Pullman Nuno Bandeira, and Yannis Papakonstantinou, investigates algorithmic issues related to

✉ Pablo Barceló
pbarcelo@uc.cl

Marco Calautti
marco.calautti@unitn.it

¹ Faculty of Mathematics, Institute for Mathematical and Computational Engineering, School of Engineering, Universidad Católica de Chile & IMFD Chile, Avda. Vicuña Mackenna 4860, Macul, 782-0436 Santiago, Chile

² Department of Information Engineering and Computer Science, University of Trento, Via Sommarive, 9 I-38123, 38123 Povo, TN, Italy

cosine similarity queries over vector databases. This problem is of practical importance, arising in important applications such as document retrieval recommender systems, and mass spectrometry. The paper deals with the efficient evaluation of such queries and provides novel optimality guarantees.

- The paper “Semi-Oblivious Chase Termination: The Sticky Case” by Marco Calautti and Andreas Pieris, studies termination of the chase, a fundamental algorithmic tool in database theory with several applications. The paper studies when the chase terminates regardless of the input data. While in general this problem is undecidable, the authors provide an elegant analysis showing that it can be solved in elementary time for the class of sticky rules, a prominent paradigm for obtaining decidability of rule-based reasoning tasks.
- The paper “Consistent Query Answering for Primary Keys in Datalog” by Paraschos Koutris and Jef Wijsen, considers the classical problem of evaluating consistent answers to conjunctive queries over databases that may violate primary key constraints. The authors show the surprising result that for any self-join free conjunctive query for which this problem can be solved in polynomial time, it can also be solved in LOGSPACE.
- Finally, the paper “On the expressive power of linear algebra on graphs”, by Floris Geerts, studies the expressive power of Matlang, a recently proposed query language for specifying both relational and linear algebra properties over matrices. This beautiful paper characterizes when Matlang can distinguish between two graphs represented by their adjacency matrices, in terms of both spectral and combinatorial properties.

Publisher’s Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.