



ELSEVIER

Contents lists available at SciVerse ScienceDirect

Journal of Computer and System Sciences

www.elsevier.com/locate/jcss


Foreword

Editor's foreword

This issue of the Journal of Computer and System Sciences contains eleven invited articles on the foundations of data management, which are representatives of the finest research in database theory of the last couple of years. These articles have been selected from the most renowned conferences in the area – the ACM Symposium on Principles of Database Systems (PODS), the International Conference on Database Theory (ICDT) and the Alberto Mendelzon Workshop on the Foundations of Data Management (AMW) – and have been thoroughly reviewed and extended following the guidelines of the Journal of Computer and System Sciences refereeing process.

This collection of paper ranges over different topics in the area of data management. In the article "On the Equivalence of Distributed Systems with Queries and Communication", Abiteboul, ten Cate and Katsis study the equivalence problem for distributed data management systems that are modeled as a collection of active XML documents. The problem is undecidable in general but they describe several relevant restrictions that lead to decidability, and even to tractability. The paper "The Language of Plain SO-tgds: Composition, Inversion and Structural Properties", by Arenas, Pérez, Riveros and Reutter, studies the feasibility of a language for describing schema mappings that is closed under composition and inverse. Although a clean language with such characteristics is unlikely to exist, they prove that under an important semantic restriction of the interpretation of mappings a natural language that meets those requirements can be found. In the paper "Solutions in XML Data Exchange", Bojanczyk, Kolodziejczyk and Murlak prove several important results for XML data exchange, most notably perhaps, that the problem of checking the existence of solutions for a given source instance can be solved in polynomial time. The article "On Simplification of Schema Mappings", by Calvanese, de Giacomo, Lenzerini and Vardi, studies the problem of whether a schema mapping can be specified in a simplified form. Previous work in the area asks for the simplified mapping to be equivalent to the original one, which is a rather restrictive requirement. The present paper studies a weaker notion based on containment instead of equivalence, and provides several techniques and results for the problem. They do this for relational and graph-based schema mappings.

In the article "Capturing Model-Based Ontology Evolution at the Instance Level: the Case of DL-Lite", Kharlamov, Zheleznyakov and Calvanese study the problem of evolution of knowledge bases in the description logic DL-lite. The focus is on identifying fragments of the logic that are closed under evolution, and studying whether the notion of evolution can be approximated in those fragments in which it cannot be expressed. In the paper "On the Aggregation Problem for Synthesized Web Services", Deng, Fan, Libkin and Wu study the aggregation problem for mediators of Web services described as deterministic finite state transducers defined in terms of templates for component services. The authors show that the decidability of the problem depends on the underlying graph of the mediator: It is undecidable for very simple acyclic graphs and decidable in single-exponential time if the underlying graph is acyclic. Restrictions that lead to tractability are also considered. The paper "Expressiveness and Static Analysis of Extended Conjunctive Regular Path Queries", by Freydenberger and Schweikardt, studies the containment and equivalence problem for an extension of the class of conjunctive regular path queries (CRPQs) with path comparisons. It also studies the succinctness and expressive power of this language in comparison with CRPQs. In the article "Simplifying XML Schema: Single-Type Approximations of Regular Tree Languages", Gelade, Idziaszek, Martens, Neven and Paredaens study the XML schema language of XSDs. In particular, how the union and difference of XSDs, which are not expressible in the language, can be approximated within the class of single-type regular tree languages. The paper provides results for both upper and lower approximations. In the paper "Semantic Query Optimization in the Presence of Types", Meier, Schmidt, Wei and Lausen present a logic-based query optimization method, with roots in the classical chase procedure, that unifies semantic and type-based optimization techniques. The method requires extending the chase to queries and dependencies that make use of negation and disjunction. In the article "A Dichotomy on the Complexity of Counting Database Repairs", Makowski and Wijsen prove a dichotomy for the complexity of the following problem: Given a database D that violates a set of primary keys K and a conjunctive query Q , how many repairs of D with respect to K satisfy Q ? It is shown that this number is either polynomially time computable or $\#P$ -complete, for the case when Q does not have self-joins. Finally, in the paper "Tractable Counting of the Answers to Conjunctive Queries", Pichler and Skritek initiate the systematic search for fragments of the class of conjunctive queries for which the problem

of counting the number of its answers is tractable. They prove that this is the case for queries of bounded hyper-treewidth and no existentially quantified variables, and that removing any of these conditions easily leads to intractability.

I am very grateful to the authors of the selected papers for their hard work in preparing the contributions, and to the reviewers for their efforts in ensuring the high quality of the present issue.

Guest Editor
Pablo Barceló
Department of Computer Science,
University of Chile,
Chile
E-mail address: pbarcelo@dcc.uchile.cl

7 December 2012
Available online 21 January 2013