Query Rewriting under Existential Rules

Andreas Pieris

School of Informatics, University of Edinburgh
apieris@inf.ed.ac.uk

Abstract of the Invited Talk

There is a clear consensus that the required level of scalability in ontology-mediated querying (OMQ) can only be achieved via query rewriting, a prominent tool that allows us to exploit standard database technology for OMQ purposes. The key idea is to reduce the problem in question to the problem of evaluating a first-order (FO) query over a relational database. This technique was originally proposed in 2005 in the context of DL-Lite. Since then it has been extensively applied, not only to more expressive DLs, but also to existential rules (a.k.a. tuple-generating dependencies and Datalog+/-rules). The goal of this invited talk is to discuss FO-rewritability of ontology-mediated queries based on the main decidable classes of existential rules, i.e., the classes of linear, (frontier-)guarded, sticky and acyclic sets of existential rules.

The first part of the talk will focus on pure FO-rewritability, where the rewriting process is database independent. For the classes of existential rules that always admit FO-rewritings, that is, linear, sticky and acyclic, I will present algorithms for constructing such rewritings, and discuss their practical relevance. For the classes that do not always admit FO-rewritings, namely (frontier-)guarded, I will discuss the challenging problem of deciding whether a rewriting exists. In view of the fact that the above (pure) FO-rewritings are unavoidably very large, the second part of the talk will focus on combined FO-rewritability, a technique that allows us to construct small rewritings at the price of touching the database (but in a controlled way).

Acknowledgements