

A Theory of Regular Queries - Abstract

Moshe Y. Vardi¹

¹University Professor and George Distinguished Service Professor in Computational Engineering at Rice University, Fellow for Science and Technology Policy at the Baker Institute for Public Policy, Houston, Texas, USA

Abstract

On 17th June 2025, Moshe Y. Vardi delivered a keynote talk at the 33rd Symposium on Advanced Database Systems in Ischia (Napoli, Italy). The following is the abstract of his talk

Abstract

A major theme in relational database theory is navigating the tradeoff between expressiveness and tractability for query languages, where the query-containment problem is considered a benchmark of tractability. The query class UCQ, consisting of unions of conjunctive queries, is a fragment of first-order logic that has a decidable query containment problem, but its expressiveness is limited. Extending UCQ with recursion yields Datalog, an expressive query language that has been studied extensively and has recently become popular in application areas such as declarative net working. Unfortunately, Datalog has an undecidable query containment problem. Identifying a fragment of Datalog that is expressive enough for applications but has a decidable query-containment problem was an open problem for several years.

In the area of graph databases, there has been a similar search for a query language that combines expressiveness and tractability. Because of the need to navigate along graph paths of unspecified length, transitive closure has been considered a fundamental operation. Query classes of increasing complexity using the operations of disjunction, conjunction, projection, and transitive closure have been studied, but the classes lacked natural closure properties. The class RQ of regular queries emerged about a decade ago as a natural query class that is closed under all of its operations and has a decidable query-containment problem.

RQ turned out to be a fragment of Datalog where recursion can be used only to express transitive closure. Further more, it turns out that applying this idea to Datalog, that is, restricting recursion to the expression of transitive closure, does yield the long-sought goal an expressive fragment of Datalog with a decidable query-optimization problem.

Declaration on Generative AI

The authors have not employed any Generative AI tools.

SEBD 2025: 33rd Symposium on Advanced Database System, June 16-19, 2025 - Ischia, Italy

 vardi_DELETE_THIS_@cs.rice.edu (M. Y. Vardi)



© 2025 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).