

Seminar Announcement

Title: Identifying the Extent of Completeness of Query Answers over Partially Complete Databases

Speaker: Simon Razniewski

Date: Friday, May 15th

Time: 15:00-15:30

Place: Seminar Room POS 1.02

Background: Joint work with Flip Korn (Google Research), Werner Nutt and Divesh Srivastava (AT&T Labs-Research), paper to be presented at SIGMOD 2015

Link to paper: https://srazniewski.files.wordpress.com/2015/05/2015_sigmod.pdf

Abstract:

In many applications including loosely coupled cloud databases, collaborative editing and network monitoring, data from multiple sources is regularly used for query answering. For reasons such as system failures, insufficient author knowledge or network issues, data may be temporarily unavailable or generally nonexistent.

Hence, not all data needed for query answering may be available. In this paper, we propose a natural class of completeness patterns, expressed by selections on database tables, to specify complete parts of database tables. We then show how to adapt the operators of relational algebra so that they manipulate these completeness patterns to compute completeness patterns pertaining to query answers.

Our proposed algebra is computationally sound and complete with respect to the information that the patterns provide. We show that stronger completeness patterns can be obtained by considering not only the schema but also the database instance and we extend the algebra to take into account this additional information. We develop novel techniques to efficiently implement the computation of completeness patterns on query answers and demonstrate their scalability on real data.