

Keynote Speech: Scalable Algorithms and Big-Data Applications by Recursive SQL Queries with Aggregates: a Datalog-enabled approach – Abstract

Carlo Zaniolo¹

¹University of California, Los Angeles, CA 90095, USA

Abstract

The use of basic SQL aggregates in recursive queries enables programmers to employ query languages to develop complete big-data applications, including graph, machine learning and data mining applications. To achieve this goal, programmers must make sure that their SQL queries can be converted into equivalent Datalog programs that combine rigorous declarative semantics with very efficient and highly scalable fixpoint based semantics. Thus, our approach provides methods and tools to verify that (i) queries with recursive aggregate have Stable Model Semantics (SMS) and (ii) such SMS can be represented via a fixpoint-based computation that is conducive to bulk-synchronous and stale-synchronous parallelism. We also provide techniques to restructure queries that satisfy (i) but not (ii) into queries that satisfy both.

Keywords

Datalog programs, Recursive SQL queries, Stable Model Semantics (SMS)

SEBD 2021: The 29th Italian Symposium on Advanced Database Systems, September 5-9, 2021, Pizzo Calabro (VV), Italy

✉ zaniolo@cs.ucla.edu (C. Zaniolo)



© 2021 Copyright for this paper by its authors. Use permitted under Creative Commons License Attribution 4.0 International (CC BY 4.0).



CEUR Workshop Proceedings (CEUR-WS.org)