

Foundations of Databases and Query Languages

Exercise 11: Datalog Evaluation

10 July 2015

Exercise 11.1 Consider the following program P :

$$\begin{aligned} T(x) &\leftarrow e(x) \\ T(x) &\leftarrow a(x, y) \wedge T(y) \wedge b(x, z) \wedge T(z) \end{aligned}$$

- (a) Describe, in your own words, the kind of structures that the query $\langle T, P \rangle$ recognises.
- (b) Give the rules for semi-naive evaluation of P .
- (c) Compute the semi-naive evaluation of P for the database that contains the following facts:

$e(1) \quad e(2) \quad e(6) \quad a(3, 1) \quad a(4, 3) \quad a(5, 3) \quad a(7, 5) \quad b(3, 2) \quad b(5, 3) \quad b(7, 6)$

Specify for each newly derived fact which of the rule(s) of (b) will produce it at the given point in the derivation.

Exercise 11.2 Consider the “Same generation” Datalog program given in the lecture (predicates S “same generation”, p “parent”, h “human”):

$$\begin{aligned} S(x, x) &\leftarrow h(x) \\ S(x, y) &\leftarrow p(x, w) \wedge S(v, w) \wedge p(y, v) \end{aligned}$$

and the adorned version for query $S(1, x)$:

$$\begin{aligned} \text{(Rule } a) \quad \text{Query}^f(x) &\leftarrow S^{bf}(1, x) \\ \text{(Rule } b) \quad S^{bf}(x, x) &\leftarrow h(x) \\ \text{(Rule } c) \quad S^{bf}(x, y) &\leftarrow p(x, w) \wedge S^{fb}(v, w) \wedge p(y, v) \\ \text{(Rule } d) \quad S^{fb}(x, x) &\leftarrow h(x) \\ \text{(Rule } e) \quad S^{fb}(x, y) &\leftarrow p(x, w) \wedge S^{fb}(v, w) \wedge p(y, v) \end{aligned}$$

together with the database that contains the following facts for predicate p :

$h(1) \quad h(2) \quad h(3) \quad h(4) \quad h(5) \quad h(6) \quad h(7)$
 $p(1, 2) \quad p(2, 3) \quad p(4, 3) \quad p(5, 4) \quad p(6, 1) \quad p(7, 1).$

Sketch the database as a tree. What are the expected answers to the query? Apply the QSQR algorithm to compute the answer to the query.

Exercise 11.3 Consider the following modified version of the same generation program:

$$\begin{aligned} S(x, x) &\leftarrow h(x) \\ S(x, y) &\leftarrow p(x, w) \wedge p(y, v) \wedge S(v, w) \end{aligned}$$

What is the adorned version of this program for query $S(1, x)$? Use this program to show that it is possible that some tuples in an input-relation are not copied to the sup_0 relation of a rule during the execution of the QSQR algorithm.

Exercise 11.4 (Abiteboul, Hull and Vianu; Exercise 13.14) Consider the following program:

$$\begin{aligned} Sv(x, y) &\leftarrow \text{flat}(x, y) \\ Sv(x, y) &\leftarrow \text{up}(x, z_1) \wedge Sv(z_1, z_2) \wedge \text{flat}(z_2, z_3) \wedge Sv(z_3, z_4) \wedge \text{down}(z_4, y) \end{aligned}$$

Give the magic set transformation for this program and query $Sv(a, y)$, where a is a constant.