


Administrivia

	Median	Max	Max Possible
MW 2	80	85	85
Q2 2	65.5	69	71
Overall	84.8	95.8	100

Final 13 Dec at 11am — 1pm

Datalog



$$P(x, y) :- E(x, y).$$

$$P(x, z) :- E(x, y), P(y, z).$$

$$P = \left\{ \begin{array}{l} (1, 2), (1, 3), (1, 4) \\ (2, 3), (2, 4) \\ (3, 4) \end{array} \right\}$$

- ① $P(x,y) :- E(x,y).$
- ② $P(x,z) :- E(x,y), P(y,z).$

Beginning of time
Iter 0

→ Bottom-up evaluation

Iter 1

Iter 2

Iter 3

Iter 4

E^0_{12}
 E^0_{23}
 E^0_{34}

E^1_{12}
 E^1_{23}
 E^1_{34}
 P^1_{12}
 P^1_{23}
 P^1_{34}

E^2_{12}
 E^2_{23}
 E^2_{34}
 P^2_{12}
 P^2_{23}
 P^2_{34}
 P^2_{13}
 P^2_{24}

E^3_{12}
 E^3_{23}
 E^3_{34}
 P^3_{12}
 P^3_{23}
 P^3_{34}

E^4_{12}
 E^4_{23}
 E^4_{34}
 P^4_{12}
 P^4_{23}
 P^4_{34}
 P^4_{13}
 P^4_{24}
 P^4_{14}

open world hypothesis
↓
closed world hypothesis

← Top-down evaluation

From this point, nothing more will ever happen.

(Under some conditions) Datalog programs must terminate.

db := input facts

Repeat forever:


db' :=

All your classical db optimizations (Indexing, pushdowns, etc.)
rules (db)

db = db ∪ rules

if db' = db: stop

else: db := db'



db = rules(db) Least fixpoint

Naive bottom-up evaluator

$$P^{(i+1)}(x, y) := P^{(i)}(x, y)$$

$$E^{(i+1)}(x, y) := E^{(i)}(x, y)$$


$$P^{(i+1)}(x, y) := E^{(i)}(x, z), P^{(i)}(z, y).$$

$$P^{(i+1)} = P^{(i)} \cup \left\{ (x, y) \mid \begin{array}{l} E^{(i)}(x, z), \\ P^{(i)}(z, y) \end{array} \right\}$$

Seminaive
evaluator

$$E^{(i)}(x, z), \Delta P^{(i)}(z, y)$$

~~$E^{(i)}(x, z), P^{(i)}(z, y)$~~



$$\Sigma^* (a+b)^* b$$

$$+ (a+b)^* b$$

The Problem of Negation

- Stratified negation
- Negation as failure