

Complexity of Evaluating Relational Queries

Given a basic graph pattern Q
& a database I

For basic graph patterns
Data complexity
— Poly($|I|$).
Query complexity
— NP-complete

checking whether Q matches I is NP-complete.

Question : Given a database I
& a RA query Q

what is the complexity of computing $Q(I)$?

— Depends on the query. ← "schema" : "Query complexity".

Some queries are easy to compute

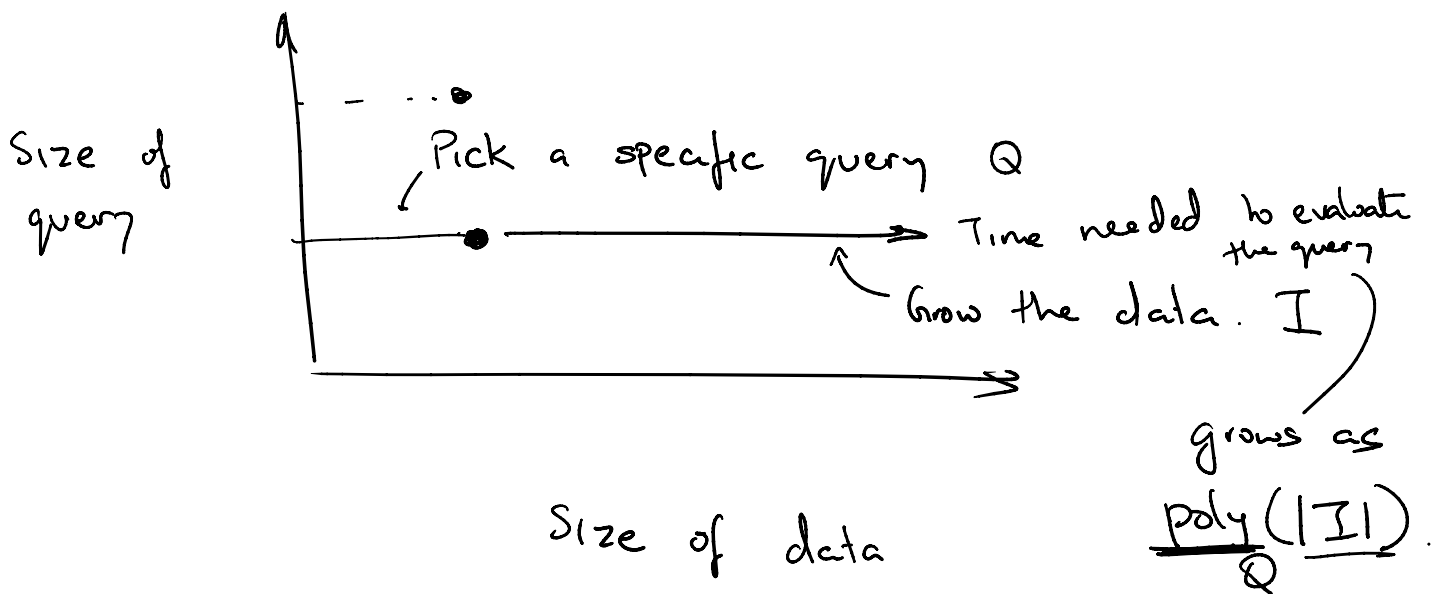
Other queries ("joins") require a lot more work.

- Depends on the data. \swarrow "schema" : "Data" complexity

Query complexity of RA. NP-complete



"Is $Q(I)$ non-empty?"

Data complexity of RA $\text{poly}(|I|)$
 AC^0



Size of query \uparrow

\uparrow Evaluate larger & larger queries \leftarrow



Normalization / Denormalization.

- └ For performance
- └ Architectural reasons

1NF
2NF
⋮
BCNF
⋮

— Moving on from graph databases,

Can we express richer forms of recursion?

- Cousins: Two people whose parents are siblings

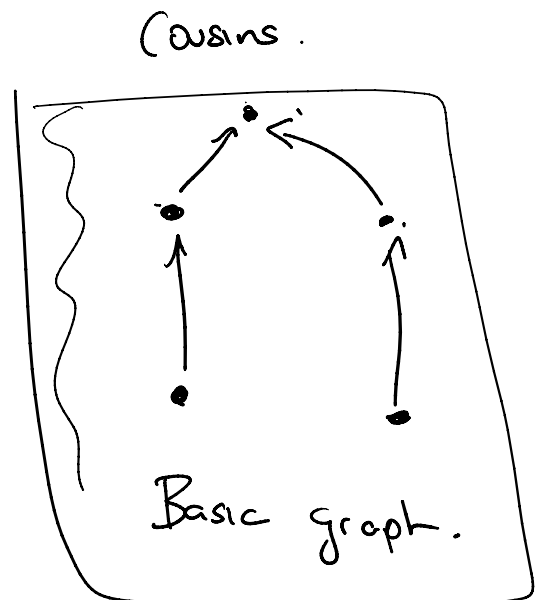
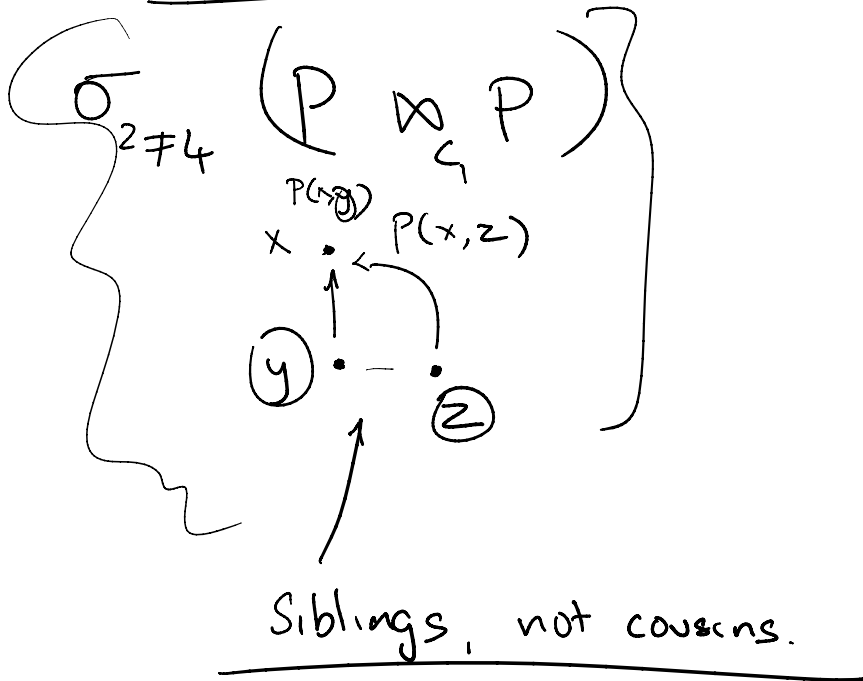
- Two people who share the same grandparent.

Parent
 $c_1 \mid c_2$

 $x \mid y$

: x is a parent of y.


How to compute cousins using RA / SQL?



$$GP = \prod_{14} (P \bowtie_{2=3} P)$$

$$Cousins = \prod_{24} (GP \bowtie_{1=3} GP)$$

4-way join in SQL/RA



For all people $x, y, z, w, u,$

If we have

Parent(x, z) and Parent(z, u)

and Parent(y, w) and Parent(w, u)

then we have

Cousins(x, y).

$$\textcircled{4} \text{ Cousins}(x, y) :- \underline{P(x, z), P(z, u), P(y, w), P(w, u)}.$$

Our first Datalog query

These commas really mean "...and..."

These commas really
mean "and".
conjunctions.

For all x, y, z, w, u , if then

Datalog "rules"

Head

Body

Cousins(x y):- $P(x$ $z)$ $P(z$ $u)$
 $P(y$ $w)$ $P(w$ $u)$

If everything in body, then head.

From body, produce head.

What might recursion be?

The body itself has head.

What about disjunctions / or s?

Just write 2/3/4/... rules.

Three "approaches" to evaluating Datalog programs

- Top down ←

- Bottom-up: Naive, semi-naive

- Magic sets ← "Query-guided evaluation"

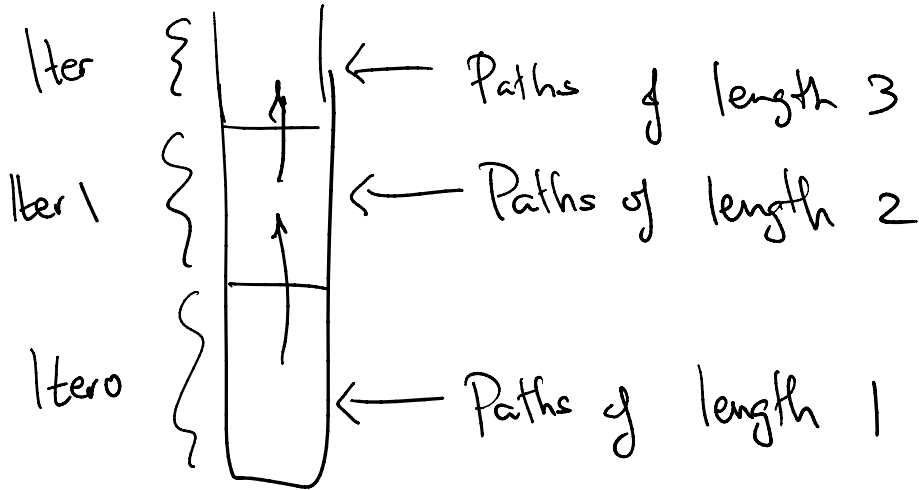
$$\forall x, y \left(\underset{\substack{\uparrow \\ \text{Contravariant}}}{x} \text{ If } \underset{\substack{\uparrow \\ \text{Covariant}}}{y} \text{ then } y \right) \iff \overset{\forall x, y}{\text{not } x} \text{ or } y.$$

$$\forall y \left(\text{not } \exists x \text{ or } y \right)$$

Theorem: Non recursive Datalog = Relation Algebra.

↑
Equal in expressive

Equal in expressive power.



But if there are n vertices in the graph, at most n^2 paths to be discovered.

"Fixed point"

It is eventually going to run out of tuples to discover