Graph Structured Data is now everywhere



- Facebook, Twitter
- DBPedia (Wikipedia represented as a graph)
- Biological databases, geological databases, social databases...

Semantic Web \longrightarrow Graph Database

Web resources are modeled as nodes of a graph, edges are relations between resources

▲□▶ ▲□▶ ▲ □▶ ▲ □▶ 亘 のへぐ

Semantic Web — Graph Database

Web resources are modeled as nodes of a graph, edges are relations between resources



Semantic Web — Graph Database

Web resources are modeled as nodes of a graph, edges are relations between resources



Querying Semantic Web:

SPARQL:

SPARQL Protocol and RDF Query Language

- Standardized query language for RDF
- Query language with strong relational flavor

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

Querying Semantic Web:

SPARQL:

SPARQL Protocol and RDF Query Language

- Standardized query language for RDF
- Query language with strong relational flavor
- Think of Relational Algebra or SQL over graphs (selection, projection, joins, unions, etc...)

(日)

Connectivity Queries: a novel challenge

Talk about paths in graphs



Connectivity Queries: a novel challenge

In a social network: Am I connected to a superstar?

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

- Workflows (biological): how is the path from process A to process B?
- Maps: shortest path form A to B?



▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

Users who might have made a typo:



▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ



▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ



▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ



▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ



▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

TriAL for RDF: Adapting Graph Query Languages for RDF Data

Leonid Libkin¹ Juan Reutter² Domagoj Vrgoč²

¹University of Edinburgh

²PUC Chile



Outline

- 1. Graph databases Most studied connectivity queries
- 2. Relationship between connectivity and recursion An algebra for querying graphs

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □

3. Where to go from here

Outline

- 1. Graph databases Most studied connectivity queries
- 2. Relationship between connectivity and recursion An algebra for querying graphs

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □

3. Where to go from here

Graph databases (for this talk):

We abstract them as triples:



Graph databases (for this talk):

We abstract them as triples:



・ロン ・聞 と ・ ヨ と ・ ヨ と

1 9 Q C

Graph databases (for this talk):

We abstract them as triples:



・ロト ・雪 ト ・ ヨ ト ・

= 900

A graph database is a collection of triples.

RDF graphs can be represented by relational databases.

RDF graphs can be represented by relational databases.



◆□ > ◆□ > ◆豆 > ◆豆 > ̄豆 _ のへで

RDF graphs can be represented by relational databases.



The distinction is in the queries that we want to ask (such as connectivity queries)



▲□▶▲圖▶▲≣▶▲≣▶ ≣ のへで



The label of the path is

revisionOf · revisionOf



The label of the path is

 $made_by \cdot worked_with \cdot worked_with$



The label of the path is

 $revisionOf^- \cdot made_by \cdot worked_with \cdot worked_with$



Revision and people involved

 $made_by \cdot (worked_with)^*$

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ



Revision and people involved

 $made_by \cdot (worked_with)^*$



Revision and people involved

 $made_by \cdot (worked_with)^*$



Revision and people involved

 $made_by \cdot (worked_with)^*$



Revision and people involved

 $made_by \cdot (worked_with)^*$



Revision and people involved

 $made_by \cdot (worked_with)^*$

RPQs is used as a primitive for querying paths

- Supported by various systems
- Simple, declarative language (easy to state what you want)

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

- Really efficient evaluation: Automata techniques allow to use fast reachability algorithms to evaluate RPQs.
- Lots of extensions

Outline

- 1. Graph databases Most studied connectivity queries
- 2. Relationship between connectivity and recursion An algebra for querying graphs

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □

3. Where to go from here

In practice, things are not that simple

Graph DB systems struggle to support RPQs.

Neo4j only supports concatenation, star

▲□▶ ▲□▶ ▲ □▶ ▲ □▶ 亘 のへぐ
In practice, things are not that simple

While RPQs are now part of SPARQL

- Semantics not clearly defined (many changes in last years)
- no clear guidelines for implementation

Problem with RPQs: limited expressivity

to study it we use SPARQL (with RPQs)

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

more or less like SQL + RPQs



▲□▶ ▲圖▶ ▲臣▶ ▲臣▶ 三臣 - 釣��



▲□▶ ▲□▶ ▲ □▶ ▲ □▶ 亘 のへぐ





Find all versions with an error

*Q*₁ : That originate from a valid version And return the latest revision



Find all versions with an error

Q₂: That originate from a valid version And the person responsible for this

What is the problem?

- Q1 can be expressed using SPARQL (essentially SQL + RPQs)
- Q₂ can NOT be expressed in SPARQL
- Conclusion:
 - Need to reason while moving along paths
 - Cant do this using RPQs or similar primitives

Problem with RPQs: implementation

Simple and efficient implementation using automata theory

▲□▶ ▲□▶ ▲ □▶ ▲ □▶ 亘 のへぐ

Problem with RPQs: implementation

Traditionally, DB systems do not implement techniques that rely on automata theory.

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

XML and XPath

Why not base implementation on relational queries

Understanding connectivity queries from a relational point of view

◆□ ▶ ◆□ ▶ ◆ □ ▶ ◆ □ ▶ ● □ ● ● ● ●

Understanding connectivity queries from a relational point of view

What are RPQs? What does (x, a^*, y) means?



► Let S ∘ S' be the composition of binary relations S and S':

 $\mathcal{S} \circ \mathcal{S}' = \{(x,z) \mid (x,y) \in \mathcal{S} \land (y,z) \in \mathcal{S}'\}$

What does (x, a^*, y) means?

Take binary relation A given by all x, y that are connected via label a in the graph.

 $(x, a^+, y) = A \cup A \circ A \cup A \circ A \circ A \cup \ldots$

Compose *A* with itself over and over again... until we reach a fixed point.

(日)







(x, revisionOf*, y)



◆□ ▶ ◆□ ▶ ◆ □ ▶ ◆ □ ▶ ◆ □ ◆ ○ へ ○

Composition is just a join!

$$S \circ S' = S \bigotimes_{2=1'}^{1,2'} S'$$

Composition is just a join!

$$S \circ S' = S \bigotimes_{2=1'}^{1,2'} S'$$

$$S(x_1, x_2) = S'(x_{1'}, x_{2'})$$

Composition is just a join!

$$S \circ S' = S \bigotimes_{2=1'}^{1,2'} S'$$

$$S(x_1, x_2) = S'(x_{1'}, x_{2'})$$

Composition is just a join!

$$S \circ S' = S \bigotimes_{2=1'}^{1,2'} S'$$

$$S(x_1, \mathbf{x}) = S'(\mathbf{x}, \mathbf{x}_{2'})$$

Composition is just a join!

$$S \circ S' = S \bigotimes_{2=1'}^{1,2'} S'$$

$$S(\mathbf{x}_1, \mathbf{x}) = S'(\mathbf{x}, \mathbf{x}_{2'})$$

Composition is just a join!

$$S \circ S' = S \bigotimes_{2=1'}^{1,2'} S'$$

$$S(x_1, x) = S'(x, x_{2'})$$

Composition is just a join!

$$oldsymbol{S} \circ oldsymbol{S}' = oldsymbol{S}^{1,2'}_{2=1'} oldsymbol{S}'$$

▲□▶ ▲□▶ ▲ □▶ ▲ □▶ 亘 のへぐ

Thus reachability is just recursive iteration of joins.

TriAL: an algebra for graphs

We now define an algebra for triples, that:

- can express RPQs
- can even express queries such as Q2

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

Is based on relational algebra

Composing ternary relation

We need to manage triples... No obvious way to do it

 $(x,y,x) \circ (x',y',z')$

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

Composing ternary relation

We need to manage triples... No obvious way to do it

$(x, y, x) \circ (x', y', z')$

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

We take the approach of relational algebra, and define all possible compositions.







$R \bowtie R'$



$R \stackrel{1,3',3}{\bowtie} R'$

◆□ > ◆□ > ◆豆 > ◆豆 > ̄豆 _ のへで



$R \overset{1,3',3}{\bowtie} R'_{2=1'}$

▲□▶ ▲□▶ ▲□▶ ▲□▶ = 三 のへで



 $E \underset{2=1'}{\overset{1,3',3}{\bowtie}} E$

▲□▶ ▲□▶ ▲ □▶ ▲ □▶ 亘 のへぐ



 $E \underset{2=1'}{\overset{1,3',3}{\bowtie}} E$



 $E \underset{2=1'}{\overset{1,3',3}{\bowtie}} E$



 $E_{2=1'}^{1,3',3}E$



 $E^{1,3',3}_{\substack{\boxtimes\\ 2=1'}}E$



 $E^{1,3',3}_{\underset{2=1'}{\bowtie}}E$
A simple join



 $E \overset{1,3',3}{\bowtie}_{2=1'} E$

▲□▶ ▲□▶ ▲ □▶ ▲ □▶ 亘 のへぐ

TriAL: An algebra of triples

R: set of triples
Relational representation of a an RDF graph

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

A TriAL expression is built using

- Set R of triples
- ► Joins 🖂
- ► Union ∪
- ► Difference \

Adding recursion – TriAL*

For binary reachability we just iterate the join

- For triples things are not symmetric
 - In particular some joins are not associative

So we need both left and right star

Adding recursion – TriAL*

For binary reachability we just iterate the join

- For triples things are not symmetric
 - In particular some joins are not associative
 - So we need both left and right star

 $(e \bowtie)^* = \emptyset \cup e \cup e \bowtie e \cup (e \bowtie e) \bowtie e \cup \dots,$

Adding recursion – TriAL*

For binary reachability we just iterate the join

- For triples things are not symmetric
 - In particular some joins are not associative
 - So we need both left and right star



Q₁: That originate from a valid version And return the latest revision

 $(E \bowtie_{3=1'}^{1,2',3'})^*$



Q₁: That originate from a valid version And return the latest revision

 $(E \bowtie_{3=1'}^{1,2',3'})^*$



Q₁: That originate from a valid version And return the latest revision

 $(E \bowtie_{3=1'}^{1,2',3'})^*$



Q₁: That originate from a valid version And return the latest revision

 $(E \bowtie_{3=1'}^{1,2',3'})^*$



Q₁: That originate from a valid version And return the latest revision

 $(E \bowtie_{3=1'}^{1,2',3'})^*$



Q₁: That originate from a valid version And return the latest revision

 $(E \bowtie_{3=1'}^{1,2',3'})^*$



Q₁: That originate from a valid version And return the latest revision

 $(E \bowtie_{3=1'}^{1,2',3'})^*$



Q₁: That originate from a valid version And return the latest revision

 $(E \bowtie_{3=1'}^{1,2',3'})^*$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \overset{1,3',3}{\bowtie})^* \overset{1,2,3'}{\underset{2=1'}{\bowtie}})^* \overset{1}{\underset{3=1',2=2'}{\bowtie}})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \overset{1,3',3}{\bowtie})^* \overset{1,2,3'}{\underset{2=1'}{\bowtie}})^* \overset{1}{\underset{3=1',2=2'}{\bowtie}})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$



*Q*₂ : That originate from a valid version And the person responsible

$$((E \bigotimes_{2=1'}^{1,3',3})^* \bigotimes_{3=1',2=2'}^{1,2,3'})^*$$

- Similar complexity bounds to RPQs
- Evaluation algorithm uses dynamic programming

▲□▶ ▲□▶ ▲ □▶ ▲ □▶ 亘 のへぐ
TriAL*: fragment of relational algebra

- well known formalism
- can be translated into SQL- like statements (or datalog-like)
- fits right onto relational query implementations but we might need new heuristics

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

Other theoretical advantages of TriAL*

know expressive power: First Order logic with transitive closure and fixed amount of variables

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

Can I pose this query?

Outline

- 1. Graph databases Most studied connectivity queries
- 2. Relationship between connectivity and recursion An algebra for querying graphs

< □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □ > < □

3. Where to go from here

What now?

Lets get back to what we know, and study connectivity queries from a relational perspective

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

- Shortest path
- Count number of paths
- Aggregate on paths (total distance, etc)

What now?

Include TriAL in graph implementations

- Include it in SPARQL
- compare performance with Neo4j, Dex

▲□▶ ▲□▶ ▲ 三▶ ▲ 三▶ 三 のへぐ

RDF DBMS (Jenna, etc)