

Introduction to property graphs: Data model and query languages


Victor MARSAULT

Université Gustave-Eiffel, CNRS, LIGM


MADICS / DOING

2022-07-11


Introduction



- DBMS = **D**ata**B**ase **M**anagement **S**ystem




- DBMS = **D**ata**B**ase **M**anagement **S**ystem




- DM = **D**ata **M**odel = *"The way data is structured"*
 - Relational ? XML ? Property graph ? RDF ? etc.

- DBMS = **D**ata**B**ase **M**anagement **S**ystem
- Query language
 - *"What can users ask for?"*



- DM = **D**ata **M**odel = *"The way data is structured"*
 - Relational ? XML ? Property graph ? RDF ? etc.

- DBMS = **D**ata**B**ase **M**anagement **S**ystem
- Query language
 - *“What can users ask for?”*



- Semantics of query
 - *“What does the query mean?”*
- DM = **D**ata **M**odel = *“The way data is structured”*
 - Relational ? XML ? Property graph ? RDF ? etc.

Example: DB for a small store

Client table

name	address
Alice	Wonderland
Bob	124 Conch St.
Charlie	1593 Broadway

Product table

name	price
Sponge	1€
Broom	5€
Rabbit	0€
Pocket Watch	100€

Example: DB for a small store

Client table

name	address
Alice	Wonderland
Bob	124 Conch St.
Charlie	1593 Broadway

Product table

name	price
Sponge	1€
Broom	5€
Rabbit	0€
Pocket Watch	100€

Order table

id	buyer	date
0	Alice	01-11-1865
1	Bob	07-07-2022

Order-content table

order_id	product
0	Rabbit
0	Pocket Watch
1	Sponge
1	Broom

Example: DB for a small store

Client table

name	address
Alice	Wonderland
Bob	124 Conch St.
Charlie	1593 Broadway

Product table

name	price
Sponge	1€
Broom	5€
Rabbit	0€
Pocket Watch	100€

Order table

id	buyer	date
0	Alice	01-11-1865
1	Bob	07-07-2022

Order-content table

order_id	product
0	Rabbit
0	Pocket Watch
1	Sponge
1	Broom

Example: DB for a small store

Client table

name	address
Alice	Wonderland
Bob	124 Conch St.
Charlie	1593 Broadway

Product table

name	price
Sponge	1€
Broom	5€
Rabbit	0€
Pocket Watch	100€

Order table

id	buyer	date
0	Alice	01-11-1865
1	Bob	07-07-2022

Order-content table

order_id	product
0	Rabbit
0	Pocket Watch
1	Sponge
1	Broom

Relational DBMS = tables with cross-references



Example: DB for a small store

Client table

name	address
Alice	Wonderland
Bob	124 Conch St.
Charlie	1593 Broadway

Product table

name	price
Sponge	1€
Broom	5€
Rabbit	0€
Pocket Watch	100€

Order table

id	buyer	date
0	Alice	01-11-1865
1	Bob	07-07-2022

Order-content table

order_id	product
0	Rabbit
0	Pocket Watch
1	Sponge
1	Broom

Relational DBMS = tables with cross-references



Example: DB for a small store

Client table

name	address
Alice	Wonderland
Bob	124 Conch St.
Charlie	1593 Broadway

Product table

name	price
Sponge	1€
Broom	5€
Rabbit	0€
Pocket Watch	100€


Order table

id	buyer	date
0	Alice	01-11-1865
1	Bob	07-07-2022

Order-content table


order_id	product
0	Rabbit
0	Pocket Watch
1	Sponge
1	Broom

Vast majority of DMBS are relational, not graph



Figure/data from db-engines.com, June 2022

Vast majority of DMBS are relational, not graph

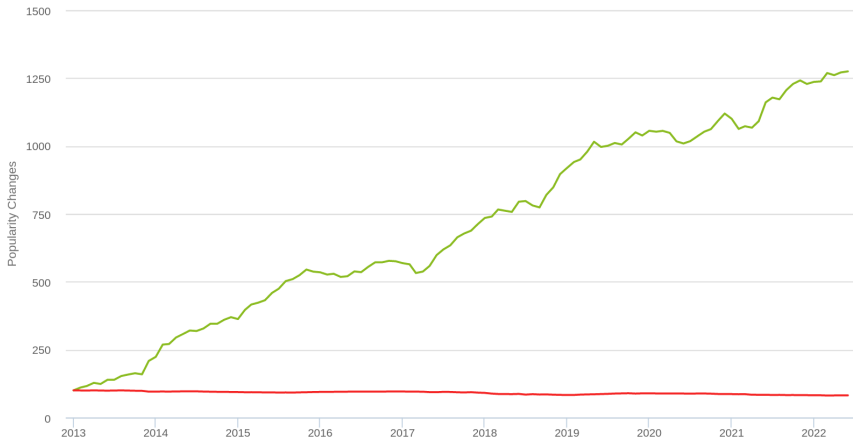


Figure/data from db-engines.com, June 2022

Graph DBMS is growing in popularity



+25% per year since 2013



Figure/data from db-engines.com, June 2022

Some data have intrinsically the structure of graphs:

- Semantic web
- Network (social, transport, etc.)
- Financial transactions

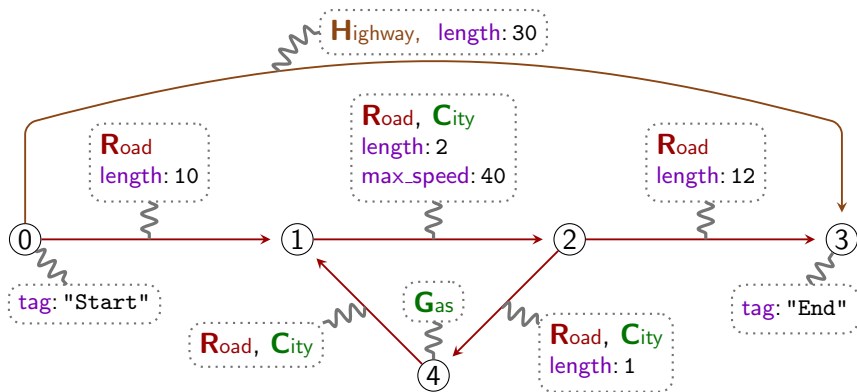
Some data have intrinsically the structure of graphs:

- Semantic web
- Network (social, transport, etc.)
- Financial transactions

Native representation of data as graphs allows:

- Specific algorithms on graphs
- Pattern matching
- More intuitive understanding/visualisation

Glimpse at our property graph example



Data model
in documentation



Property graph
(vendor dependent)

DBMS



Neo4j,
TigerGraph,
etc.

Abstract
data model



Abstract
property graph

Data model
in documentation




Property graph
(vendor dependent)

DBMS



Neo4j,
TigerGraph,
etc.

Dataless model



Graph

Abstract
data model



Abstract
property graph

Data model
in documentation



Property graph
(vendor dependent)

DBMS



Neo4j,
TigerGraph,
etc.

Part 1: without data

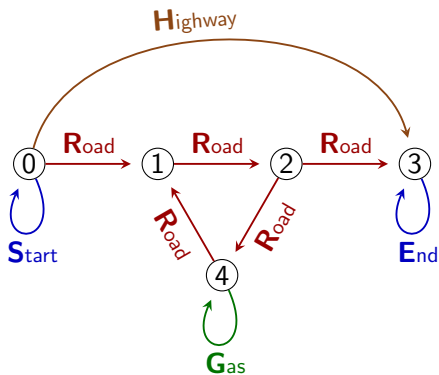
- Graphs as database
- Foundation of querying graph databases: RPQs
- Ensuring finiteness

Part 2: with data

- Property graphs
- Cypher
- GQL
- Appendix

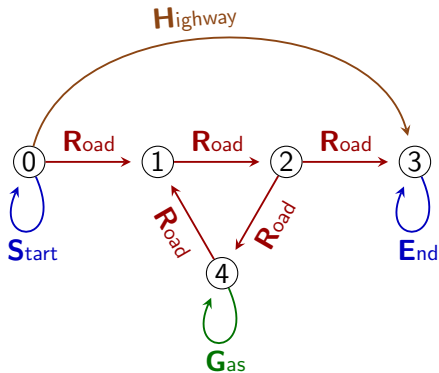
Graphs as database

A graph consists of ...




A graph consists of ...

- Vertices (a.k.a. Nodes)
0, 1, 2, 3, 4




A graph consists of ...

- Vertices (a.k.a. Nodes)
0, 1, 2, 3, 4
- Edges (a.k.a. Relationships)
 $0 \rightarrow 0$, $0 \rightarrow 1$, etc



A graph consists of ...

- Vertices (a.k.a. Nodes)
0, 1, 2, 3, 4
- Edges (a.k.a. Relationships)
 $0 \rightarrow 0$, $0 \rightarrow 1$, etc
- Edge labels: {R, H, G, S, E}




A graph consists of ...

- Vertices (a.k.a. Nodes)
0, 1, 2, 3, 4
- Edges (a.k.a. Relationships)
 $0 \rightarrow 0$, $0 \rightarrow 1$, etc
- Edge labels: {R, H, G, S, E}

Walk

- a.k.a. Path
- Sequence of edges
- Can reuse vertices and edges




A graph consists of ...

- Vertices (a.k.a. Nodes)
0, 1, 2, 3, 4
- Edges (a.k.a. Relationships)
0 → 0, 0 → 1, etc
- Edge labels: {R, H, G, S, E}

Walk

- a.k.a. Path
- Sequence of edges
- Can reuse vertices and edges

0 → 1 → 2 → 4




A graph consists of ...

- Vertices (a.k.a. Nodes)
0, 1, 2, 3, 4
- Edges (a.k.a. Relationships)
0 → 0, 0 → 1, etc
- Edge labels: {R, H, G, S, E}

Walk

- a.k.a. Path
- Sequence of edges
- Can reuse vertices and edges

0 → 1 → 2 → 4 → 4




A graph consists of ...

- Vertices (a.k.a. Nodes)
0, 1, 2, 3, 4
- Edges (a.k.a. Relationships)
 $0 \rightarrow 0$, $0 \rightarrow 1$, etc
- Edge labels: {R, H, G, S, E}

Walk

- a.k.a. Path
- Sequence of edges
- Can reuse vertices and edges

$0 \rightarrow 1 \rightarrow 2 \rightarrow 4 \rightarrow 1 \rightarrow 2 \rightarrow 3$



A graph can be stored in tables


Table for nodes

id
0
1
2
⋮

Table for **R**_{oad} edges

source	target
0	1
1	2
⋮	⋮

One table for **S**_{tart},
One table for **H**_{ighway},
etc...



A graph can be stored in tables


Table for nodes

id
0
1
2
⋮

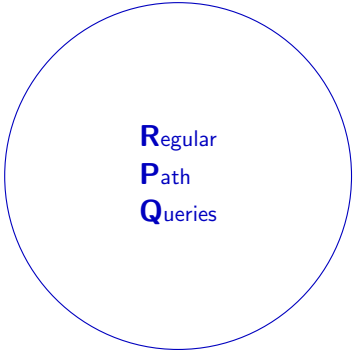
Table for **R**_{oad} edges

source	target
0	1
1	2
⋮	⋮

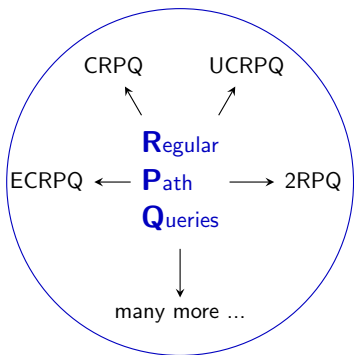
One table for **S**_{tart},
One table for **H**_{ighway},
etc...

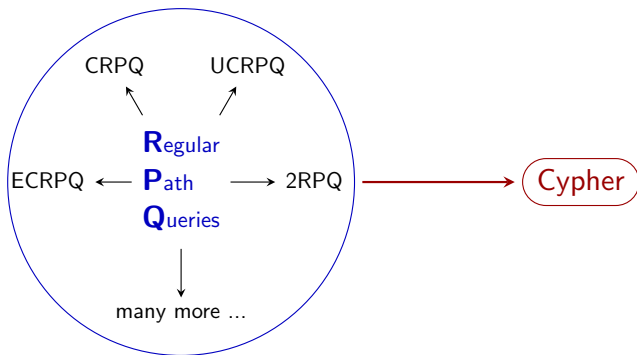


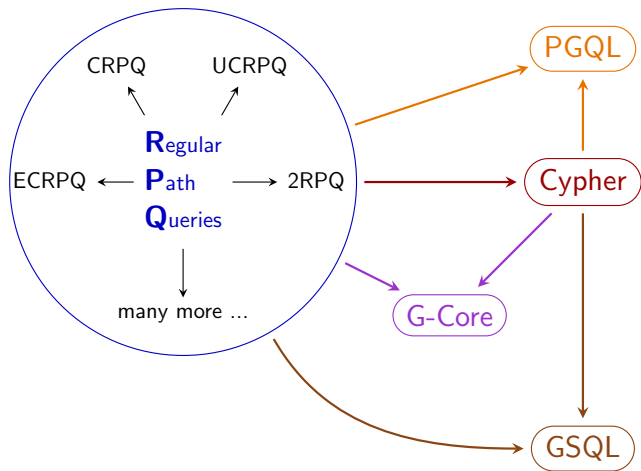
Foundation of querying graph databases: RPQs

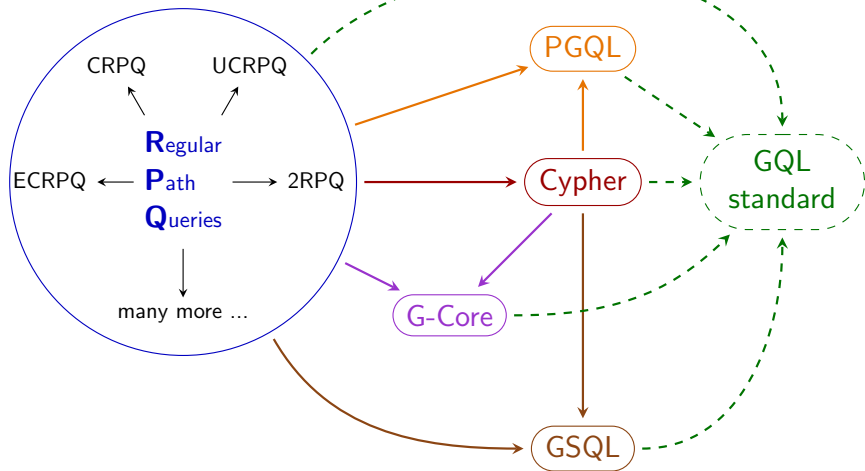


Regular
Path
Queries









$Q ::=$	A	Atoms
	QQ	Concatenation
	$Q + Q$	Union/Disjunction
	Q^*	Kleene star

where **A** is a label in the graph.

An RPQ denotes a set of words

S (**H** + **R**)^{*} **E** denotes the words of the shape **S** <something> **E**

Any number of **H** and **R**, in any order

$Q ::=$	A	Atoms
	QQ	Concatenation
	$Q + Q$	Union/Disjunction
	Q^*	Kleene star

where **A** is a label in the graph.

An RPQ denotes a set of words

S (**H** + **R**)^{*} **E** denotes the words of the shape **S** <something> **E**


Any number of **H** and **R**, in any order

Matches

A match for Q is any walk w such that Q denotes the label of w

Query **R** matches...

...every **R**_{oad}-edge



Query **R** matches...

...every **R**_{oad}-edge


0 → 1

2 → 4

1 → 2

4 → 1

2 → 3



Query **R** matches...

...every **R**_{oad}-edge

0 → 1

2 → 4


1 → 2

4 → 1

2 → 3

Query **RR** matches...

...walks of two **R**_{oad}-edges



Query **R** matches...

...every **R**_{oad}-edge

0 → 1

2 → 4

1 → 2

4 → 1

2 → 3

Query **RR** matches...

...walks of two **R**_{oad}-edges


0 → 1 → 2

2 → 4 → 1

1 → 2 → 3

4 → 1 → 2

1 → 2 → 3



Query **R** matches...

...every **R**_{oad}-edge

0 → 1

2 → 4

1 → 2

4 → 1

2 → 3

Query **RR** matches...

...walks of two **R**_{oad}-edges

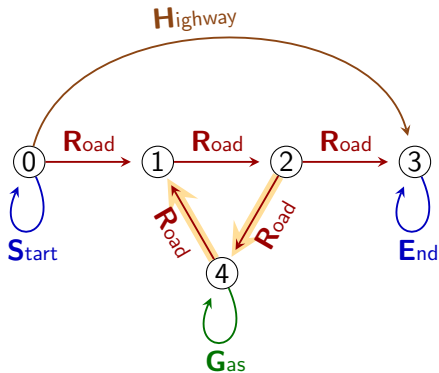
0 → 1 → 2

2 → 4 → 1

1 → 2 → 3

4 → 1 → 2

1 → 2 → 3



Query **R+H** matches...

...edges labelled by **R** or by **H**

0 → 1


2 → 4

1 → 2

4 → 1

2 → 3

0 → 3



Query **R+H** matches...

...edges labelled by **R** or by **H**

0 → 1


1 → 2

2 → 3

2 → 4

4 → 1

0 → 3



Query **R+H** matches...

...edges labelled by **R** or by **H**

0 → 1


2 → 4

1 → 2

4 → 1

2 → 3

0 → 3



Query **R+H** matches...

...edges labelled by **R** or by **H**

$0 \rightarrow 1$

$2 \rightarrow 4$

$1 \rightarrow 2$

$4 \rightarrow 1$


$2 \rightarrow 3$

$0 \rightarrow 3$

Query **S(H+R)** matches

$0 \rightarrow 0 \rightarrow 1$

$0 \rightarrow 0 \rightarrow 3$



Query **R+H** matches...

...edges labelled by **R** or by **H**

$0 \rightarrow 1$

$1 \rightarrow 2$

$2 \rightarrow 3$

$2 \rightarrow 4$


$4 \rightarrow 1$

$0 \rightarrow 3$

Query **S(H+R)** matches

$0 \rightarrow 0 \rightarrow 1$

$0 \rightarrow 0 \rightarrow 3$



Query **R+H** matches...

...edges labelled by **R** or by **H**

0 → 1

2 → 4

1 → 2

4 → 1


2 → 3

0 → 3

Query **S(H+R)** matches

0 → 0 → 1

0 → 0 → 3




Query R^* matches...

...walks consisting of R_{road} -edges

```

    0 → 1
    0 → 1 → 2
    0 → 1 → 2 → 3
    0 → 1 → 2 → 4
      ⋮
    1 → 2 → 4 → 1
      ⋮
    1 → 2 → 4 → 1 → 2 → 4 → 1
      ⋮
  
```



Query R^* matches...

...walks consisting of R_{road} -edges

$0 \rightarrow 1$

$0 \rightarrow 1 \rightarrow 2$

$0 \rightarrow 1 \rightarrow 2 \rightarrow 3$

$0 \rightarrow 1 \rightarrow 2 \rightarrow 4$


\vdots

$1 \rightarrow 2 \rightarrow 4 \rightarrow 1$

\vdots

$1 \rightarrow 2 \rightarrow 4 \rightarrow 1 \rightarrow 2 \rightarrow 4 \rightarrow 1$

\vdots



Query R^* matches...

...walks consisting of R_{road} -edges

$0 \rightarrow 1$

$0 \rightarrow 1 \rightarrow 2$

$0 \rightarrow 1 \rightarrow 2 \rightarrow 3$

$0 \rightarrow 1 \rightarrow 2 \rightarrow 4$


\vdots

$1 \rightarrow 2 \rightarrow 4 \rightarrow 1$

\vdots

$1 \rightarrow 2 \rightarrow 4 \rightarrow 1 \rightarrow 2 \rightarrow 4 \rightarrow 1$

\vdots



Query R^* matches...

...walks consisting of R_{road} -edges

$0 \rightarrow 1$

$0 \rightarrow 1 \rightarrow 2$

$0 \rightarrow 1 \rightarrow 2 \rightarrow 3$

$0 \rightarrow 1 \rightarrow 2 \rightarrow 4$


\vdots

$1 \rightarrow 2 \rightarrow 4 \rightarrow 1$

\vdots


$1 \rightarrow 2 \rightarrow 4 \rightarrow 1 \rightarrow 2 \rightarrow 4 \rightarrow 1$

\vdots




$$Q_1 = S(R+H)^*E$$

Q_1 matches...



$$Q_1 = \mathbf{S(R+H)^*E}$$

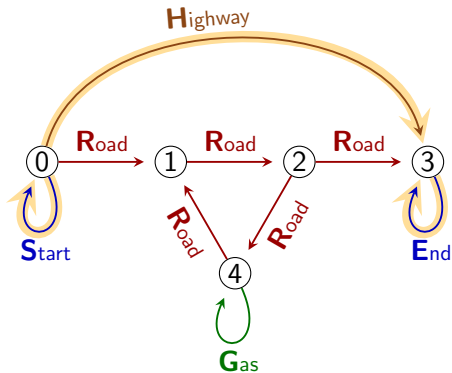
Q_1 matches...



$$Q_1 = S(R+H)^*E$$

Q_1 matches...


- The highway



$$Q_1 = S(R+H)^*E$$

Q_1 matches...


- The highway
- The direct road



$$Q_1 = S(R+H)^*E$$

Q_1 matches...

- The highway
- The direct road
- Road with laps in the circuit




$$Q_1 = S(R+H)^*E$$

Q_1 matches...

- The highway
- The direct road
- Road with laps in the circuit

$$Q_2 = S(R+H)^*G(R+H)^*E$$

Q_2 matches...




$$Q_1 = S(R+H)^*E$$

Q_1 matches...

- The highway
- The direct road
- Road with laps in the circuit

$$Q_2 = S(R+H)^*G(R+H)^*E$$

Q_2 matches...



$$Q_1 = S(R+H)^*E$$


Q_1 matches...


- The highway
- The direct road
- Road with laps in the circuit

$$Q_2 = S(R+H)^*G(R+H)^*E$$

Q_2 matches...

- Road with laps in the circuit






⚠ Answer may be infinite ⚠

Ensuring finitess

Definition

- Returns the endpoints of matches




Definition

- Returns the endpoints of matches

$$Q_1 = \mathbf{S} (\mathbf{R} + \mathbf{H})^* \mathbf{E}$$

$$Q_2 = \mathbf{S} (\mathbf{R} + \mathbf{H})^* \mathbf{G} (\mathbf{R} + \mathbf{H})^* \mathbf{E}$$

- All matches are: $0 \rightarrow \dots \rightarrow 3$
 $\Rightarrow Q_1$ and Q_2 return $\{(0, 3)\}$



Pros and cons

Pros

- Efficient algorithms
- Well grounded theory

Pros and cons

Pros

- Efficient algorithms
- Well grounded theory


Cons

- Limited information in the answer
 - User: *"I want to go from Paris to Lyon by car"*
 - Database: *"Yes you can"*

PGQL (Oracle), GSQL (TigerGraph), G-core [Angles et al. 2018]

Definition

- Return the walk with the least number of edges




PGQL (Oracle), GSQL (TigerGraph), G-core [Angles et al. 2018]

Definition

- Return the walk with the least number of edges

$$Q_1 = S (R+H)^* E$$

- Q_1 returns 1 walks
 - the highway
- Walks taking the road have more edges



PGQL (Oracle), GSQL (TigerGraph), G-core [Angles et al. 2018]

Definition


- Return the walk with the least number of edges

$$Q_1 = S (R+H)^* E$$

- Q_1 returns 1 walks
 - the highway
- Walks taking the road have more edges

$$Q_2 = S (R+H)^* G (R+H)^* E$$

- Q_1 returns 1 walks
 - the one-lap road



Pros

- Returns walks
- Efficient algorithms
- Horizontal post-processing
 - Horizontal = along the walk
 - *“Is there a gas station on the way?”*
 - *“What is the length of the walk?”*

Pros and cons

Pros

- Returns walks
- Efficient algorithms
- Horizontal post-processing
 - Horizontal = along the walk
 - *“Is there a gas station on the way?”*
 - *“What is the length of the walk?”*


Cons

- No vertical post-processing
 - Vertical = accross the walks with the same endpoints
 - *“What is the shortest walk in time?”*
 - *“What is the connectedness level?”*

Cypher (Neo4j)

Definition

- Return walks
- Forbidden to repeat edges




Cypher (Neo4j)

Definition

- Return walks
- Forbidden to repeat edges

$$Q_1 = S (R+H)^* E$$

- Q_1 returns 2 walks
 - the highway
 - the straight road




Cypher (Neo4j)

Definition

- Return walks
- Forbidden to repeat edges

$$Q_1 = S (R+H)^* E$$

- Q_1 returns 2 walks
 - the highway
 - the straight road
- Walks with circuit laps
 - \Rightarrow repeat the middle edge



Cypher (Neo4j)

Definition


- Return walks
- Forbidden to repeat edges

$$Q_1 = S (R+H)^* E$$

- Q_1 returns 2 walks
 - the highway
 - the straight road
- Walks with circuit laps
 - \Rightarrow repeat the middle edge

$$Q_2 = S (R+H)^* G (R+H)^* E$$

- Q_2 returns nothing



Pros and cons

Pros

- Returns walks
- Counting matches
- Horizontal and vertical post-processing

Pros and cons

Pros

- Returns walks
- Counting matches
- Horizontal and vertical post-processing


Cons

- All problems are computationally hard [Martens et al. 2020]
 - Counting, enumeration, existence
 - Checking whether Q_2 returns anything → Already hard
- No repeated edge → not always intuitive

New theoretical proposal [David-Francis-Marsault 2023?]

Definition

- Returns walks
- Each edge may match each atom only once




New theoretical proposal [David-Francis-Marsault 2023?]

Definition

- Returns walks
- Each edge may match each atom only once

$$Q_2 = S (R+H)^* G (R+H)^* E$$




New theoretical proposal [David-Francis-Marsault 2023?]

Definition

- Returns walks
- Each edge may match each atom only once

$$Q_2 = S (R+H)^* G (R+H)^* E$$

- Returns the 1-lap road only




New theoretical proposal [David-Francis-Marsault 2023?]

Definition

- Returns walks
- Each edge may match each atom only once

$$Q_2 = S (R+H)^* G (R+H)^* E$$

- Returns the 1-lap road only
 - Before **G** → use the left **R**




New theoretical proposal [David-Francis-Marsault 2023?]

Definition

- Returns walks
- Each edge may match each atom only once

$$Q_2 = S (R+H)^* G (R+H)^* E$$

- Returns the 1-lap road only
 - Before **G** → use the left **R**
 - After **G** → use the right **R**




New theoretical proposal [David-Francis-Marsault 2023?]

Definition

- Returns walks
- Each edge may match each atom only once

$$Q_2 = \mathbf{S} (\mathbf{R} + \mathbf{H})^* \mathbf{G} (\mathbf{R} + \mathbf{H})^* \mathbf{E}$$

- Returns the 1-lap road only
 - Before **G** → use the left **R**
 - After **G** → use the right **R**
- > 1 circuit lap ⇒ some edge use the same atom twice




New theoretical proposal [David-Francis-Marsault 2023?]

Definition

- Returns walks
- Each edge may match each atom only once

$$Q_2 = \mathbf{S} (\mathbf{R} + \mathbf{H})^* \mathbf{G} (\mathbf{R} + \mathbf{H})^* \mathbf{E}$$

- Returns the 1-lap road only
 - Before **G** → use the left **R**
 - After **G** → use the right **R**
- > 1 circuit lap \Rightarrow some edge use the same atom twice



$$Q_1 = \mathbf{S} (\mathbf{R} + \mathbf{H})^* \mathbf{E}$$

- Returns the highway and the straight road

Pros and cons

Pros

- Returns walks
- Horizontal and vertical post-processing
- Enumerating matches is efficient
- Counting matches


Pros and cons

Pros

- Returns walks
- Horizontal and vertical post-processing
- Enumerating matches is efficient
- Counting matches

Cons


- Counting matches is computationally hard
- Answer depends on the way the query is written
 - R^* allows no lap in the circuit
 - R^*R^* allows 1 lap in the circuit
 - $R^*R^*R^*$ allows 2 lap in the circuit
 - etc.



⚠ Answer may be infinite ⚠

- Several way to ensure finiteness
 - Homomorphism semantics → Returns little information
 - Shortest-walk semantics → No vertical post-processing
 - Trail semantics → Not efficient in bad cases
 - etc.
- No solution is clearly superior

Property graphs




Vertices and edges may bear:

- zero or more labels
- zero or more properties

- Property = key-value pair
- Key = string
- Value = bool, int, str, ...

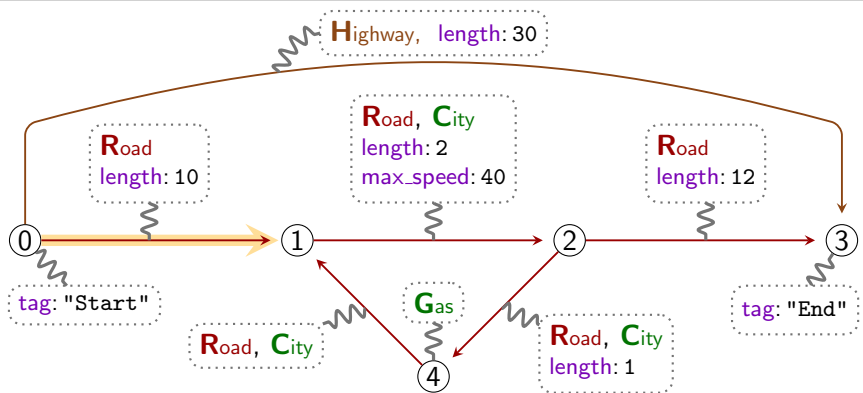
Storing our property graph into tables



id	source_id	target_id	Road	Highway	City	length	max_speed
e ₀₁	0	1	true	false	false	10	
e ₁₂	1	2	true	false	true	10	40
e ₄₁	4	1	true	false	true		
⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮

GraphLabelsProperties

Storing our property graph into tables




	id	source_id	target_id	Road	Highway	City	length	max_speed	
▶	e ₀₁	0	1	true	false	false	10		◀
	e ₁₂	1	2	true	false	true	10	40	
	e ₄₁	4	1	true	false	true			
	⋮	⋮	⋮	⋮	⋮	⋮	⋮	⋮	

GraphLabelsProperties

Wedding is edgy...


id	person1	person2
e_{12}	Alice	Bob
\vdots	\vdots	\vdots



..but trouble is trouble


Wedding is edgy...

id	person1	person2
e_{12}	Alice	Bob
\vdots	\vdots	\vdots



..but trouble is trouble

id	person1	person2	person3
e_{123}	Alice	Bob	Eve
\vdots	\vdots	\vdots	\vdots



Cypher

History

- Created for Neo4j DBMS in 2011
- Neo4j and Cypher become popular
- Project openCypher since 2015
 - Goal: standardize Cypher
 - Evolution led by community (oCIM)
 - Formal semantics [Francis et al. 2018] [Green et al. 2019]
 - Replaced by GQL
- Implemented in SAP HANA, Redis graph, etc.

- Vertices: `MATCH (:Gas)`

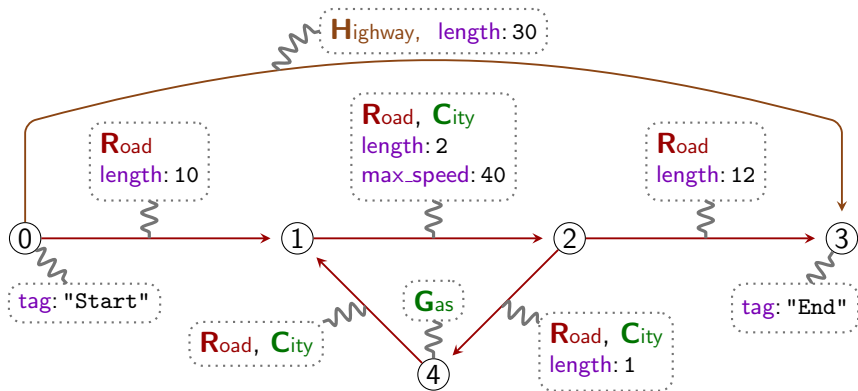
- Vertices: `MATCH (:Gas)`
- Edges: `MATCH -[:Road]->`

- Vertices: `MATCH (:Gas)`
- Edges: `MATCH -[:Road]->`
- Disjunction: `MATCH -[:Road|Highway]->`

- Vertices: `MATCH (:Gas)`
- Edges: `MATCH -[:Road]->`
- Disjunction: `MATCH -[:Road|Highway]->`
- Concatenation: `MATCH ()-[:Road]->(:Gas)-[:Road]->()`

- Vertices: `MATCH (:Gas)`
- Edges: `MATCH -[:Road]->`
- Disjunction: `MATCH -[:Road|Highway]->`
- Concatenation: `MATCH ()-[:Road]->(:Gas)-[:Road]->()`
- Kleene star: `MATCH ()-[:Road*]->()`

- Vertices: `MATCH (:Gas)`
- Edges: `MATCH -[:Road]->`
- Disjunction: `MATCH -[:Road|Highway]->`
- Concatenation: `MATCH ()-[:Road]->(:Gas)-[:Road]->()`
- Kleene star: `MATCH ()-[:Road*]->()`
- Q_1 :
`MATCH ({tag:"Start"})-[:Road|Highway*]->({tag:"End"})`




Query

```
MATCH (s)-[:City]->(t)
```

Result

s	t
1	2
2	4
4	1




Query

```
MATCH (s)-[:City]->(t)
```

Result

s	t
1	2
2	4
4	1

Cypher returns a table...




Query

```
MATCH (s)-[:City]->(t)
```

Result

s	t
1	2
2	4
4	1




Query

```
MATCH (s)-[:City]->(t)
```

Result

s	t
1	2
2	4
4	1

Cypher returns a table... but computes walks



Query


MATCH

```
(s {name:"Start"})  
  -[:Road|Highway*]->  
    (t {name:"End"})
```

Result

s	t
0	3
0	3

Cypher returns a table... but computes walks



Query

MATCH

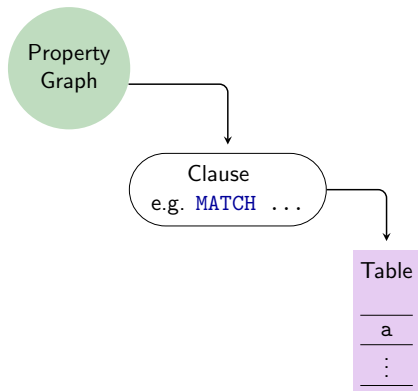
```
(s {name:"Start"})  
  -[:Road|Highway*]->  
    (t {name:"End"})
```

Result

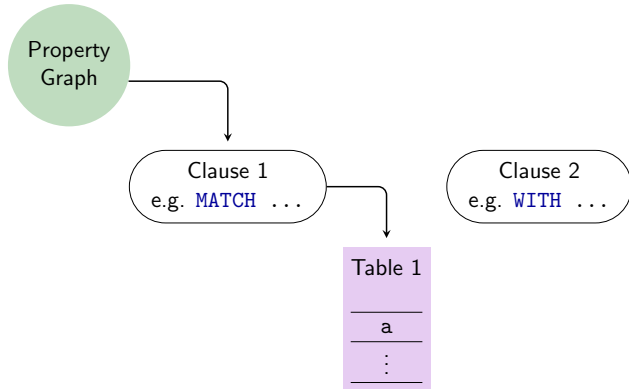
s	t	
0	3	← The highway
0	3	← The direct road

- **ORDER BY:** orders row
- **WHERE:** filters row
- **WITH** or **RETURN:**
 - adds/renames columns
 - horizontal aggregation: e.g. reduce
 - vertical aggregation: e.g. count
- **CREATE/DELETE/SET:** updates the property graph


A Cypher query actually chain clauses




A Cypher query actually chain clauses




A Cypher query actually chain clauses




A Cypher query actually chain clauses





Example

- Clause 1 makes some pattern matching
- Clause 2 aggregates over the result of Clause 1



Example

- Clause 1 makes some pattern matching
- Clause 2 aggregates over the result of Clause 1

⇒ Trail semantics (rich post-processing at the cost of efficiency)

GQL

Since 2015: openCypher aims at making Cypher a standard

Since 2015: openCypher aims at making Cypher a standard

2017:

- Cypher won't become a standard !
- Merge existing languages instead ?

Since 2015: openCypher aims at making Cypher a standard

2017:

- Cypher won't become a standard !
- Merge existing languages instead ?

2018: Comparison of features in Cypher, G-Core, PGQL

Since 2015: openCypher aims at making Cypher a standard

2017:

- Cypher won't become a standard !
- Merge existing languages instead ?

2018: Comparison of features in Cypher, G-Core, PGQL

2019: ISO project [ISO/IEC CD 39075 and 9075-16.2]

- Academia
- Neo4j (Cypher)
- Oracle (SQL, PGQL)
- TigerGraph (GSQL)
- ...

Since 2015: openCypher aims at making Cypher a standard

2017:

- Cypher won't become a standard !
- Merge existing languages instead ?

2018: Comparison of features in Cypher, G-Core, PGQL

2019: ISO project [ISO/IEC CD 39075 and 9075-16.2]

- Academia
- Neo4j (Cypher)
- Oracle (SQL, PGQL)
- TigerGraph (GSQL)
- ...

2021: Pattern matching is decided [Alin et al. 2022]

Since 2015: openCypher aims at making Cypher a standard

2017:

- Cypher won't become a standard !
- Merge existing languages instead ?

2018: Comparison of features in Cypher, G-Core, PGQL

2019: ISO project [ISO/IEC CD 39075 and 9075-16.2]

- Academia
- Neo4j (Cypher)
- Oracle (SQL, PGQL)
- TigerGraph (GSQL)
- ...

2021: Pattern matching is decided [Alin et al. 2022]

2023 (expected): Publication of version 1 of GQL Standard

An RPQ may have infinitely many matches


- GQL has to ensure finiteness of answer
- No solution is clearly superior

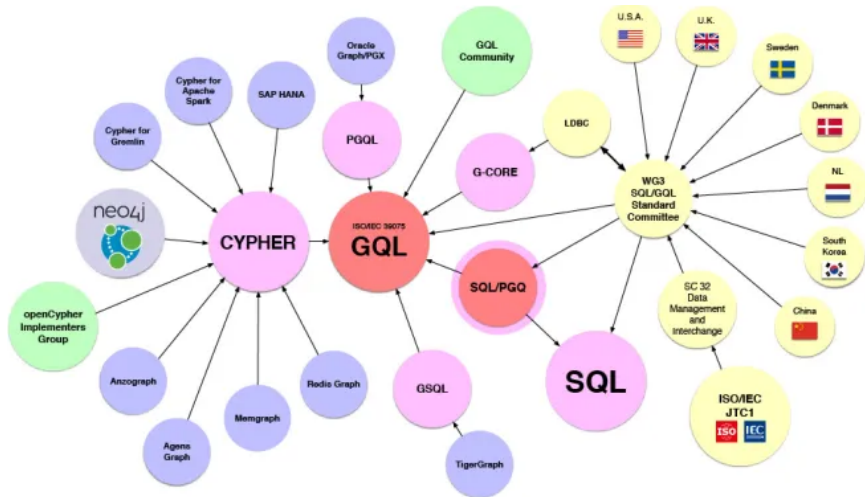
An RPQ may have infinitely many matches

- GQL has to ensure finiteness of answer
- No solution is clearly superior

GQL does not choose

- Trail semantics → TRAIL
- Shortest-walk semantics → SHORTEST
- Syntax restriction → WALK
- ...





Thank you for your attention !

- Introduction
 - General setting 1
 - Relational DBMS 2
 - Graph DBMS in practice 3
 - Abstraction 7

- Graphs as database
 - Definition by example 9
 - Graph vs Tables 10

- Foundation of querying graph databases: RPQs
 - RPQ = Regular expression 12
 - Examples 13
 - Main queries 16
 - The challenge with RPQs 17


- Ensuring finiteness
 - Homomorphism semantics 18
 - Shortest walk 20
 - Trail semantics 22
 - Run-based semantics 24


- Property graphs
 - Definition by example 27
 - Property Graph vs Tables 28

- Cypher
 - History 30
 - Syntax 31
 - Examples 32

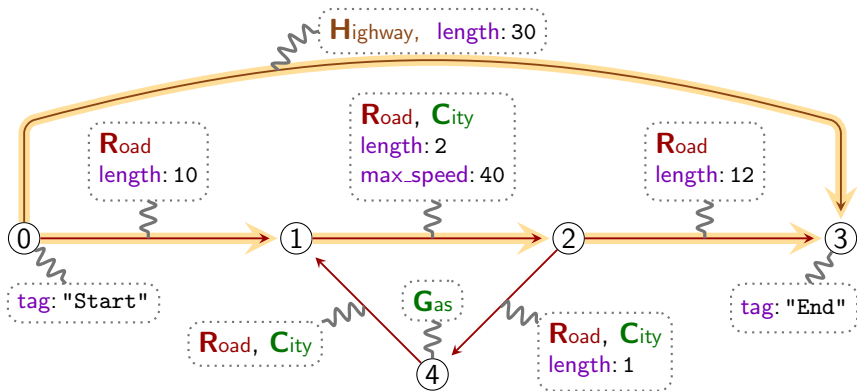
- GQL
 - History 40
 - GQL semantics 36

- Appendix






Output of GQL: set of path bindings
Path binding = walk annotated with variables



```
MATCH TRAIL (a WHERE a.tag="Start")
  [ -[r:Road]-> | -[c:City]-> ]* (b WHERE b.tag="End")
```



```
MATCH TRAIL (a WHERE a.tag="Start")
  [ -[r:Road]-> | -[c:City]-> ]* (b WHERE b.tag="End")
```

```
0 → 1 → 2 → 3
a r  r  r b
```

```
0 → 1 → 2 → 3
a r  c  r b
```