# LDBC Graph Query Language Task Force 

Status Report - 9th $^{\text {th }}$ TUC Meeting

Speaker: Hannes Voigt, TU Dresden

## Aim

- Study query languages for graph data management systems, specifically systems storing "Property Graph" data
- Query language should cover the needs of important use cases: social network benchmark, interactive and BI workloads
- Members
- Renzo Angles, Universidad de Talca
- Marcelo Arenas, PUC Chile - task force lead
- Pablo Barceló, Universidad de Chile
- Peter Boncz, Vrije Universiteit Amsterdam
- George Fletcher, Eindhoven University of Technology
- Claudio Gutierrez, Universidad de Chile
- Tobias Lindaaker, Neo Technology
- Marcus Paradies, SAP
- Raquel Pau, UPC
- Arnau Prat, UPC / Sparsity
- Juan Sequeda, Capsenta
- Oskar van Rest, Oracle Labs
- Hannes Voigt, TU Dresden
- Yinglong Xia, IBM


## Motivation

- Currently Babel of graph QLs with diversity in syntax and semantics
- PGQL: iso/homomorphism
- Cypher: "edge" isomorphism
- Gremlin: homomorphism
- SPARQL: homomorphism
- Reachability queries vs. path queries
- Practical consequences
- Applications are not portable
- Hard to define benchmarks
- Hard to compare Graph DBMS
- Standard

- Prevents vendor lock-in
- Fosters true performance competition $\rightarrow$ improvement of systems


## Closed Query Languages



HEAVEN


## Cross the Concept Chasm with Composability

- Users talk about...
- Application entities
- e.g. discussions, topics, communities, etc.
- Likely multiple abstraction levels
- Base data contains...
- Fine granular data
- Low abstraction
- E.g. individual twitter messages, retweet relationships, etc.



## High-level Design Goals

- Query Language for Property Graph Model
- Power comparable to SQL 92

- Composability (language closed over data model)
- Orthogonal language concepts
- Paths as first class citizen

| Logical Methods in Computer Science |  |  |
| :--- | :--- | :--- |
| Vol. $11(4: 1) 2015$, pp. $1-39$ | Submitted | Apr. 8, 2014 |
| www.Imcs-online.org | Published | Oct. |

EXPRESSIVE PATH QUERIES ON GRAPHS WITH DATA*
PABLO BARCELÓ ${ }^{a}$, GAELLE FONTAINE ${ }^{b}$, AND ANTHONY WIDJAJA LIN ${ }^{c}$

## Where are we now?

## Catalog of Desired Query Functionalities

- Adjacency Queries
- Graph Pattern Matching
- Navigational Queries
- Aggregate Queries
- Sub queries



## Adjacency queries

- Property access
- Get the firstName and lastName of a person having email "\$email"
- Neighborhood of a node
- Get the firstName and lastName of the friends of a person identified by email "\$email".
- K-neighborhood of a node
- Get the email, firstName and lastName of friends of the friends of a person having email "\$email" (excluding the start person) (i.e. get a list of recommended friends) (directed 2-neighborhood)


## Graph Pattern Matching

- Join
- Get the creationDate and content of the messages created by a person identified by email "\$email1" and commented by another person identified by "\$email2".
- Union
- Get the creationDate and content of the messages either created or liked by a person identified by email "\$email".
- Intersection
- Get the email, firstName and lastName of the common friends between two persons identified by emails "\$email1" and "\$email2" respectively.


## - Difference

- Given two friends identified by emails "\$email1" and "\$email2" respectively, get the email, firstName and lastName of the friends of the second person which are not friends of the first person (this questions is relevant for friendship recommendations).
- Optional
- Given a person identified by email "\$email", get the title of all the messages created by such person, and the content of the first comment replying each message (if it exists).
- Filter
- Get the properties of the people whose firstName includes the string "xxx" (it implies use of wildcards).


## Navigational queries

- Reachability
- Is there a friendship connection between two persons identified by emails "\$email1" and "\$email2" respectively?
- All Path Finding
- Get the friendship paths between two persons identified by emails "\$email1" and "\$email2" respectively.
- Shortest Path Finding
- The shortest friendship path between two persons identified by emails "\$email1" and "\$email2" respectively".
- Regular Path Query
- Get the firstName of friends of the friends of the friends of a person identified by email "\$email".
- Conjunctive Regular Path Queries
- Given a target message created on "\$dateTime" by a person identified by email "\$email", for each comment replying the target message, get the comment's content and the email of the comment's creator.
- Filtered regular path query
- Given a person identified by email "\$email", get the title of all the messages liked by such person between "\$dateTime1" and "\$dateTime2".


## Data Model



- Lis an infinite set of (node and edge) labels;
- K is an infinite set of property names
- $\mathbf{V}$ is an infinite set of literals (actual values);
- $\mathbf{T}$ is a finite set of value types (INT, VARCHAR, etc.)
- $G$ is a finite set of graphs;
- $N$ is a finite set of nodes;
- $E$ is a finite set of edges such that $N$ and $E$ have no elements in common;
- $\rho: E \rightarrow(N \times N)$ is a total function;
- $\lambda:(N \cup E) \rightarrow \operatorname{SET}(\mathrm{L})$ is a total function;
- $\sigma:(N \cup E) \times \mathbf{K} \rightarrow \operatorname{SET}(\mathbf{V})$ is a partial function;
- $\vartheta: \mathbf{V} \rightarrow \mathbf{T}$ is a function;

Single Graph example

- $\mathbf{L}=\{$ Person, knows $\}$
- $\mathbf{P}=\{i d$, name, since $\}$
- $V=\{" 123$ ", " 456 ", "Juan Sequeda", "Marcelo Arenas", "2010"\}
- $N=\{n 1, n 2\}$
- $\mathrm{E}=\{\mathrm{e} 1\}$
- $\boldsymbol{\rho}=[\rho(e 1)=(n 1, n 2)]$
- $\lambda_{\text {"knows" }}=[\lambda(n 1)=$ "Person", $\lambda(n 2)=$ "Person", $\lambda(e 1)=$
- $\boldsymbol{\sigma}=[$ $\sigma(n 1)=\{($ "id"", "123"), ("name"," "Juan Sequeda") $\}$, $\sigma(n 2)=\{(" i d "$, " "456"),' ("name"," "Marcelo
Arenas" $)\left\{\left({ }^{2}\right.\right.$ "since", " 2010 " $\left.)\right\}$


## Graphs as Tables

- Complete the picture; integrates different perspectives
- Allows to define semantics based on well-defined relational semantics
- Helps to better see/understand the delta to SQL
- Helps integration/adaption of QL in relational systems
- Suggests one possible implementation


## Graphs as Tables

- built-in properties id\$, src\$, dest\$, graph\$;
- id\$: $(G \cup N \cup E) \rightarrow$ OID;
- $\operatorname{OID}(X)=\{x . i d \$ \mid x \in X\}$
- $\operatorname{src\$ :~} E \rightarrow$ OID(N);
- dest\$: $E \rightarrow$ OID(N);
- graph\$: $N \rightarrow$ OID(G).
- two tables Vertice and Edges:
- the schema of Vertice is \{id\$, graph\$\} $U\{k \mid \sigma(n, k)$ is defined; $n \in N$ and $k \in \mathbf{K}\}$
- the schema of Edges is $\{\mathbf{s r c} \boldsymbol{\$}, \operatorname{dest} \boldsymbol{\$}\} \cup\{k \mid \sigma(e, k)$ is defined; $e \in E$ and $k \in \mathbf{K}\}$
- Vertice $=U_{n \in N} \pi_{s(\text { Vertice })}\{n\}$ Edge $=U_{n \in N} \pi_{\text {S(Edge) }}\{n\}$


## Example

Data graph G


## Query <br> SELECT x , y <br> FROM G (x:green)-[e:+]->(y:green)



## Dealing with Objects

- Nodes and edges are object with a system managed identity
- Object Immutability
- For queries, objects (nodes and edges) are immutable
- Query can create new (transient) objects out of queried data
- Consequence: Within scope of a query the object identity functionally determines meta type, label, and property values of an object
- Identity Generation
- Object constructor produces new object identities (OID values)
- The scope of ID uniqueness is the transaction (query)
- Repeatability of ID generation is not guarantied


## Current Discussion Points

## How to represent Paths

- How to represent path in the data model?
- Just use existing elements of the data model
- Keeps data model simple, but complicate interpretation on top
- As a data type for properties
- Introduced non-atomic type to properties $\rightarrow$ complicates language
- By-elements (actual node and edges) vs. by-reference (list of ids)
- Current favorite: Logical paths
- Another top level set in the data model
- Does not contain all paths but just marks paths of users interest


## Example: Logical Paths

Data graph


- $G=(N, E, P, \rho, \lambda, \delta, \sigma)$
- $N=\{n 1, \mathrm{n} 2, \mathrm{n} 3, \mathrm{n} 4, \mathrm{n} 5\}$
- $E=\{\mathrm{e} 1, \mathrm{e} 2, \mathrm{e} 3, \mathrm{e} 4, \mathrm{e} 5\}$
- $P=\varnothing$
- $\rho(\mathrm{e} 1)=(\mathrm{n} 1, \mathrm{n} 2), \rho(\mathrm{e} 2)=(\mathrm{n} 2, \mathrm{n} 3)$, $\rho(\mathrm{e} 3)=(\mathrm{n} 3, \mathrm{n} 2), \rho(\mathrm{e} 4)=(\mathrm{n} 2, \mathrm{n} 4)$, $\rho(\mathrm{e} 5)=(\mathrm{n} 3, \mathrm{n} 5)$

Result graph
SELECT p
FROM G p=((x:green)-[e:*]->(y:green)) e3


- $G^{\prime}=\left(N^{\prime}, E^{\prime}, P^{\prime}, \rho^{\prime}, \lambda^{\prime}, \delta^{\prime}, \sigma^{\prime}\right)$
- $N^{\prime}=\{n 1, n 2, n 3, n 4\}$
- $E^{\prime}=\{\mathrm{e} 1, \mathrm{e} 2, \mathrm{e} 3, \mathrm{e} 4\}$
- $P^{\prime}=\{p 1, p 2, \ldots\}$
- $\rho^{\prime}(e 1)=(n 1, n 2), \rho^{\prime}(e 2)=(n 2, n 3)$, $\rho^{\prime}(\mathrm{e} 3)=(\mathrm{n} 3, \mathrm{n} 2), \rho^{\prime}(\mathrm{e} 4)=(\mathrm{n} 2, \mathrm{n} 4)$
- $\delta^{\prime}(p 1)=[e 1, e 4]$
$\delta^{\prime}(\mathrm{p} 2)=[\mathrm{e} 1, \mathrm{e} 2, \mathrm{e} 3, \mathrm{e} 4]$


## Projection

- Relational-like projection
- Limit result of single query to single type of nodes and out edge
- SELECT X.name, X.gender | length(p) AS sim FROM G (X:Person)-[p:friend+]->(Y:Person) GRAPH BY Y
- UNION allows to assemble more complex graphs
- Unclear how edge targets are projected if of another node type
- Graph transformation-like projection (graph projection for short)
- Allows to project to multiple node and edge types in one go
- Result specified by a subgraph pattern
- Intuitive, very graphy, not 100\% orthogonal to UNION


## Example: Graph Projection

## Data graph G



Creates new edge in

```
Query:
SELECT (x)-[:{l=length(p)}]->(y)
FROM G (x:gree)-[CHEAPEST p:*]->(y:blue)
```



| id | label | E.id | E.dest | E.l |
| :--- | :--- | :--- | :--- | :--- |
| 1 | green | 21 | 4 | 4 |
|  |  | 23 | 6 | 3 |
| 6 | blue | no out edges |  |  |
| 4 | blue | no out edges |  |  |

## Example: Graph Projection w/ Aggregation

## Data graph G



## Query:

SELECT (x)-[:\{l=AVG(length(p)) $]$ ]->(z GROUP BY x:blue) FROM G (x:green)-[CHEAPEST p:*]->(y:blue)


| id | label | E.id | E.dest | E.l |
| :--- | :--- | :--- | :--- | :--- |
| 1 | green | 21 | 7 | 3,5 |
| 7 | blue | no out edges |  |  |

## Summary

- Accomplished
- Catalog of functionalities
- Core data model
- Principles object identities
- In discussion
- Data model extension for path representation
- Principles of graph construction
- Ahead of us
- Putting pieces together - define semantics of language core
- Define a syntax
- Extend core toward advanced concepts (e.g. path with data)

