Traversing and Querying Graphs with Gremlin and PGQL on Oracle Spatial and Graph

Sungpack Hong

Research Director

Oracle Labs

Martin Sevenich

Principal Researcher

Oracle Labs

Oskar van Rest

Principal Member of Technical Staff Oracle Spatial and Graph

Martijn Dwars

Senior Researcher

Oracle Labs

Safe Harbor

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Statements in this presentation relating to Oracle's future plans, expectations, beliefs, intentions and prospects are "forward-looking statements" and are subject to material risks and uncertainties. A detailed discussion of these factors and other risks that affect our business is contained in Oracle's Securities and Exchange Commission (SEC) filings, including our most recent reports on Form 10-K and Form 10-Q under the heading "Risk Factors." These filings are available on the SEC's website or on Oracle's website at http://www.oracle.com/investor. All information in this presentation is current as of September 2019 and Oracle undertakes no duty to update any statement in light of new information or future events.

Abstract

• PGQL and Gremlin are two open source languages for querying and traversing property graphs. Although the two languages look very different on the surface, there are many parallels. This session introduces both languages, shows how common queries can be expressed in either language, and shows how to use both languages with Oracle Spatial and Graph. Finally, it shows how users can compile many of their existing Gremlin queries into PGQL queries.

*Please note that this session includes research and development efforts that are not shipping product features. Oracle's Safe Harbor provisions apply to the contents of this presentation.

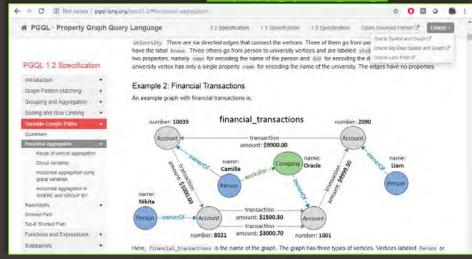


Property Graph Query Language (PGQL)

- PGQL is a "graph pattern matching" query language
 - Define a query as a high-level pattern and match it against graph data
- PGQL is aligned to SQL where possible
 - Provides familiarity for SQL users
- PGQL is part of Oracle [Big Data] Spatial and Graph products

```
/* PGQL: find top-10 people with highest number of friends */
SELECT n.firstName, COUNT(*) AS friendCount
   FROM socialNetwork
   MATCH (p:Person) -[:likes]-> (friend:Person)
GROUP BY p
ORDER BY friendCount DESC
LIMIT 10
```

PGQL home: http://pgql-lang.org/



Oracle Spatial and Graph home:

https://www.oracle.com/database/technologies/spatialandgraph.html

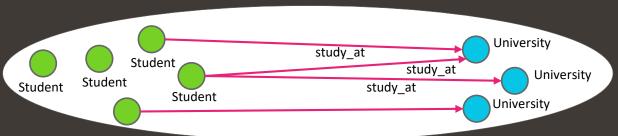


Property Graphs (1/2)

- A property graph consists of:
 - A set of vertices, each having:
 - A set of properties (key-value pairs)
 - A set of **labels** (character strings)
 - A set of edges, each having:
 - A source vertex and destination vertex
 - A set of **properties**
 - A set of labels
- Graphs are often aligned to the mental model that humans have of their data
 - Graphs make it **easy to analyze data**, even for non-programmers
 - Graphs can be **visualized naturally**, as graphs

Example: "student data" as **graph**

student network (graph)



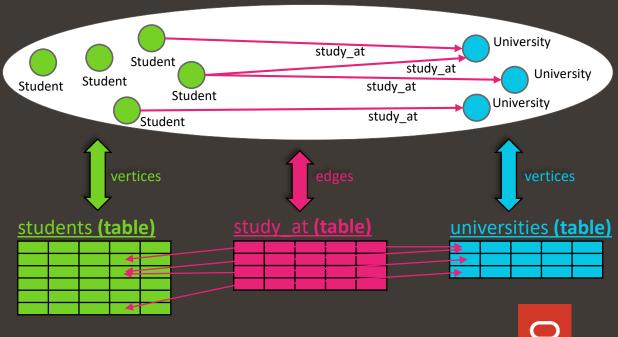


Property Graphs (2/2)

- Graphs abstract over how data is stored
 - Graph can be stored as set of tables,
 XML or JSON document, set of triples,
 adjacency lists (an efficient graph format), etc.
- Example graphs:
 - Bill of materials (product hierarchy)
 - Financial networks
 - Social networks
 - Physical networks
 - Knowledge graphs

Example: "student data" as **graph** vs. **set of tables**

student network (graph)

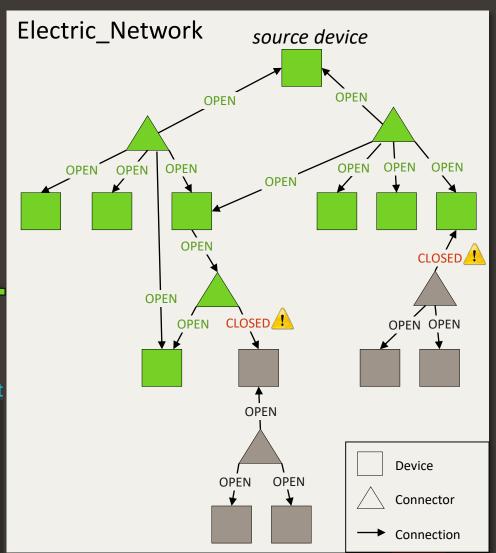


Example: impact of network disruption

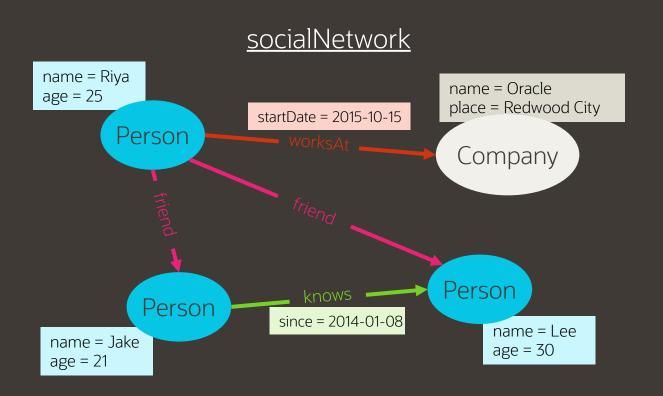


A power outage caused connectivity failures in an electric network (CLOSED). Can devices still reach each other via alternative routes?

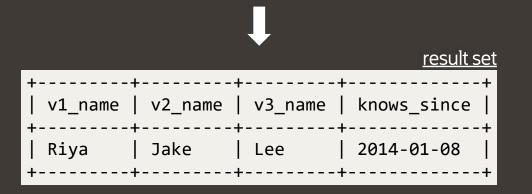
```
PATH connects to
           AS (from) \leftarrow [c1] - (connector) - [c2] \rightarrow (to)
           WHERE c1.status = 'OPEN'
                                                                          property
             AND c2.status = 'OPEN'
                                                                            graph
  SELECT n.nickname, COUNT(m)
     FROM Electric Network
   MATCH (n:Device) -/:connects to*/-> (m:Device)
   WHERE java_regexp_like(n.nickname, 'Regulator')
      AND n <> m
                                                                          result set
GROUP BY n
                                                 n.nickname
ORDER BY COUNT(m) DESC, n.nickname
                                                 Regulator, VREG2 A
                                                                         1596
                                                 Regulator, VREG4 B
                                                                         1537
                                                 Regulator, VREG4 C
                                                                         1537
                                                 Regulator, HVMV_Sub_RegA
                                                 Regulator, HVMV Sub RegB
      Copyright © 2019 Oracle and/or its affiliates.
```



Graph pattern matching



PGQL query: "find all people known by friends of Riya"





Graph queries vs. tabular queries

- PGQL queries are expressed very succinctly by means of path patterns
 - A path pattern consists of vertex patterns (e.g. (n:Person)) and edge patterns (e.g. -[e:likes]->)
- Compared to joins in SQL, path patterns increase the level of abstraction
 - Single edge pattern may correspond to multiple joins (between multiple pairs of tables)
 - Single edge pattern may correspond to recursive join (e.g. * for zero or more matches)
 - Etc.

```
WITH temp(device_id, device_name) AS (
  -- Anchor member:
 SELECT device id, name
  FROM Devices
 WHERE name = 'Regulator, HVMV_Sub_RegB'
UNI ON ALL
  -- Recursive member:
  SELECT Devices. device id, Devices. name
  FROM temp, Devices, Connections conn1,
        Connections conn2, Connectors
  WHERE temp. device_id = conn1. to_device_id
   AND conn1. from_connector_i d = Connectors. connector_i d
   AND Connectors. connector id = conn2. from connector id
   AND conn2. to device id = Devices. device id
   AND temp. device_id ! = Devices. device_id)
CYCLE device_id SET cycle TO 1 DEFAULT 0
SELECT DISTINCT device_name
FROM temp
WHERE cycle = 0
 AND device_name <> 'Regulator, HVMV_Sub_RegB'
```

```
PATH connects_to AS (from) <- (connector) -> (to)

SELECT y. name

MATCH (x: Device) -/: connects_to*/-> (y: Device)

WHERE x. name = 'Regulator, HVMV_Sub_RegB'

AND x <> y

PGOL query against a property graph
```

equivalent

Example **SQL** query vs. **PGQL** query

SQL query against a set of tables



PGQL and SQL

PGQL uses syntax and semantics of standard SQL where possible:

- Same syntactic query structure: SELECT ... FROM ... WHERE ... GROUP BY ... HAVING ... ORDER BY ...
- Same expressions and aggregations:
 IS [NOT] NULL, CAST, CASE, IN, MIN/MAX/AVG/SUM, etc.
- Same (tabular) subqueries:
 Existential (EXISTS) subqueries, scalar subqueries, and more coming
- Same data types and expressions on those data types:
 DATE, TIME [WITH TIME ZONE], TIMESTAMP [WITH TIME ZONE], etc.

 EXTRACT(MONTH FROM my_date), etc.



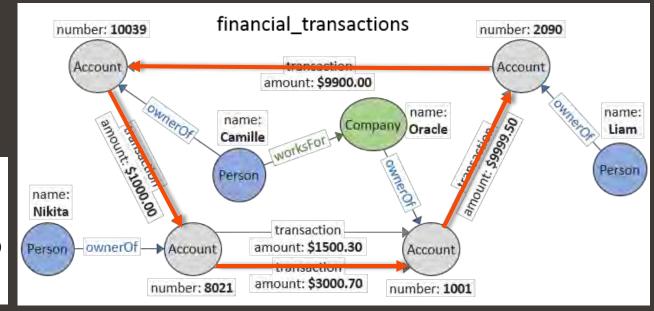
Example: financial fraud detection

Find suspicious circular payments using a **SHORTEST** path query

Note: a payment is circular when it starts and ends in the same entity

PGQL query

Property graph





PGQL in Oracle Spatial and Graph (OSG)

Graph Query (PGQL)

```
/* find friends of friends of Clara */
SELECT fof.name
  FROM myGraph
MATCH (p:Person) -/:knows{2}/-> (fof:Person)
WHERE p.name = 'Clara'
```

PGQL on In-memory Analyst (PGX)

- Excels in computationally intense workloads and recursive gueries
- Can combine graph algorithms with graph queries

Analytical graph query

PGQL on RDBMS

- Excels in workloads with mixtures of read and write queries
- Can query data sets that don't fit into the memory of a single machine

Transactional graph query

In-memory Analyst (PGX)



Bulk Update

- Synchronizes an in-memory graph snapshot with graph changes from RDBMS
- Every x seconds/minutes/hours or upon request



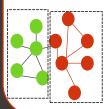


Interaction with graph algorithms

Oracle Spatial and Graph provides a rich set of built-in graph algorithms

(as part of the In-memory Analyst / PGX)

Detecting Components and Communities



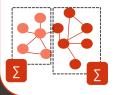
Tarjan's, Kosaraju's,
Weakly Connected
Components, Label
Propagation (w/ variants),
Soman and Narang's
Sparcification

Ranking and Walking



Pagerank, Personalized Pagerank, Betweenness Centrality (w/ variants), Closeness Centrality, Degree Centrality, Eigenvector Centrality, HITS, Random walking and sampling (w/ variants)

Evaluating Community Structures



Conductance, Modularity Clustering Coefficient (Triangle Counting) Adamic-Adar

Link Prediction

SALSA (Twitter's Who-to-follow)

Path-Finding



Hop-Distance (BFS) Dijkstra's, Bi-directional Dijkstra's Bellman-Ford's

Other Classics

Vertex Cover Minimum Spanning-Tree (Prim's)

- PGQL graph queries are often used in combination with built-in graph algorithms
 - Example:
 - Run Pagerank algorithm
 → each vertex gets a new "pagerank" property
 - 2. Run PGQL query to find the top-10 vertices with highest pageRank (ORDER BY Vertex.pagerank DESC)

Upcoming: PGX Algorithm for defining your own graph algorithms





What is Gremlin?

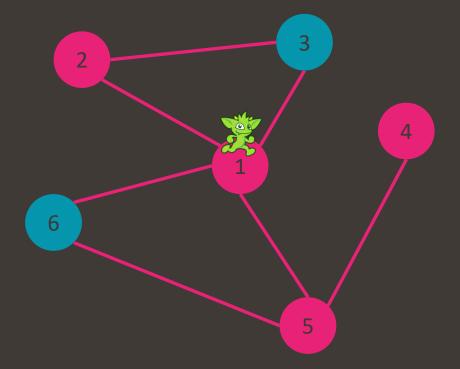
- Gremlin is a functional data-flow language
- Gremlin is part of Apache TinkerPop



- Gremlin allows to express traversals on graphs
- Each traversal consists of a sequence of (nested) steps
 - A step performs an atomic operation on the data-stream

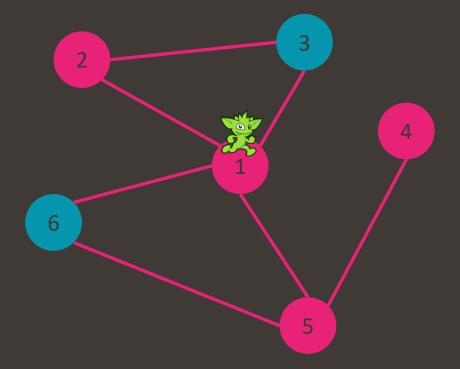
```
g.V(1)
.out('red')
.has('color', 'blue')
.count()
```

```
g.V(1)
  .out('red')
  .has('color', 'blue')
  .count()
```



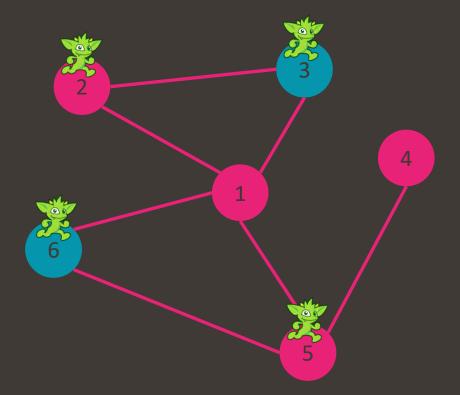


```
g.V(1)
  .out('red')
  .has('color', 'blue')
  .count()
```



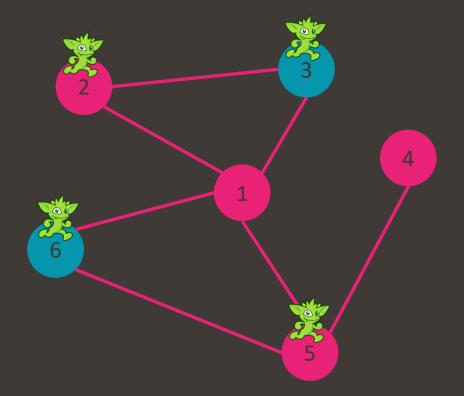


```
g.V(1)
  .out('red')
  .has('color', 'blue')
  .count()
```



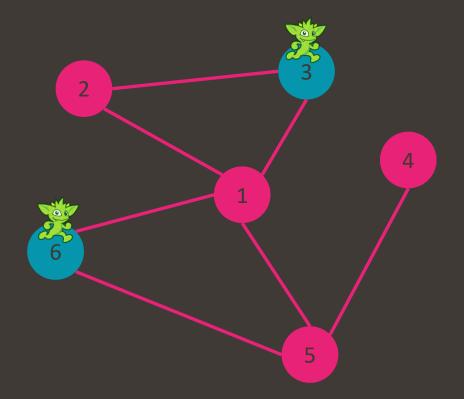


```
g.V(1)
  .out('red')
  .has('color', 'blue')
  .count()
```



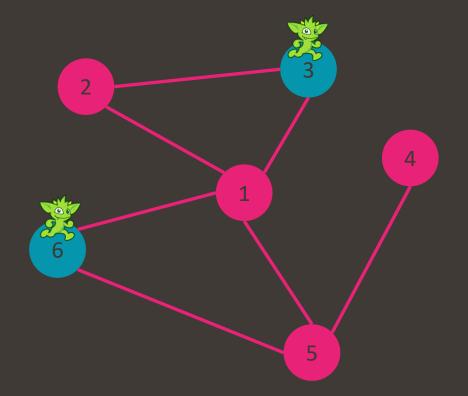


```
g.V(1)
  .out('red')
  .has('color', 'blue')
  .count()
```





```
g.V(1)
  .out('red')
  .has('color', 'blue')
  .count()
> 2
```





Gremlin Steps

Gremlin has a large variety of different steps

- Transform
 - Neighbors
 - Map to property value
- Filter
- Aggregations
- Side-effects
- Branch
 - If-then-else
 - Loops



PGOL vs. Gremlin

- PGQL is SQL-like
 - Familiar syntax & semantics
 - Declarative
- Query optimizer
 - PGQL can use its sophisticated query optimizer
 - No need to implement separate Gremlin optimizer / execution engine
- PGQL is easier to maintain



PGQL vs. Gremlin

Gremlin

```
g.V()
 .hasLabel('Customer').as('c')
                                       users have to always keep
 .values('name').as('name')
                                         the current context in
 .outE('Bought').as('b')
                                               mind
 .values('discount').as(')
 .select('b')
 .values('quantity').as('quant')
 .select('b')
 .outV()
 .has('Product', 'category', eq('toy')).as('p')
 .values('price').as('price')
 .select('price', 'disc', 'quant')
 .map{it.get().price*(1-it.get().disc/100.0)*it.get().quant}
   .as('sale price')
 .select('name','sale_price')
 .group().by('name').by(sum())
```

PGOL vs. Gremlin

PGQL

```
SELECT
    c.name, sum(p.price*(1-b.discount/100.0)*b.quantity)
MATCH
    (c:Customer) -[b:Bought]-> (p:Product)
WHERE
    p.category = 'toy'
GROUP BY
    c.name
```





From Gremlin to PGQL

- Experimental feature developed by Oracle Labs
- Automatically convert a Gremlin traversal to a PGQL query
- Users can write Gremlin but use the PGQL execution engine

From Gremlin to PGQL

- Motivation:
 - No need to reinvent the wheel
 - Support Gremlin and PGQL with only the PGQL execution engine
 - Focus on improving one engine instead of two
 - Gremlin can benefit from existing PGQL query optimizer
 - Help users transition from Gremlin to PGQL
 - keep existing Gremlin traversals
 - Implement new queries in PGQL
 - Execute both on the same engine
 - Provide tool to auto-translate existing traversals if needed



Workflow

Gremlin Query

```
g.out()
.distinct()
```



Workflow

Gremlin Query

```
g.out()
  .distinct()
```

Gremlin AST

```
[GraphStep(vertex,[]
),
VertexStep(OUT,verte
x),
DedupGlobalStep()]
```



Workflow

Gremlin Query

```
g.out()
  .distinct()
```

Gremlin AST

```
[GraphStep(vertex,[]
),
VertexStep(OUT,verte
x),
DedupGlobalStep()]
```

PGQL AST

```
Query(SelectClause(
    Distinct(),
    rRef("v2")

Attern(Vertices(["
    Transformation rules
```



Workflow

Gremlin Query

```
g.out()
  .distinct()
```

Gremlin AST

```
[GraphStep(vertex,[]
),
VertexStep(OUT,verte
x),
DedupGlobalStep()]
```

PGQL AST

```
Query(SelectClause(
    Distinct(),
VarRef("v2")
    ),
GraphPattern(Vertices(["
v1", "v2"]),
[Edge("v1")])
)
```

PGQL Query

```
SELECT DISTINCT
(v2)
MATCH
(v1) -> (v2)
```



Transformation Rules

- Each Gremlin step has a transformation rule to PGQL
- The tool keeps track of the current context of the traversal
 - Each step might alter the context
 - Which alters the SELECT of the resulting query



Gremlin to PGQL: Transformations

Gremlin

```
g.V()
    .as("a")
    has("name", "gremlin")
    out("created")
    in("created")
    where(neq("a"))
    groupCount().by("title")
```

PGQL

```
SELECT

a
MATCH

(a) [:created] > (b)
WHERE

a.name = "gremlin"
GROUP

BY

C.title
```



Gremlin to PGQL: Transformations

Gremlin

```
g.V()
    .as("a")
    .has("name", "gremlin")
    out("created")
    in("created")
    where(neq("a"))
    groupCount().by("title")
```

PGQL

```
SELECT

a
MATCH

(a) [:created] > (b)
WHERE

a.name = "gremlin"

GROUP

BY

Cotitle
```



Gremlin

The query result changes with the context

g.V() .as("a") .has("name", "gremlin") .out("created") .in("created") where(neq("a")) groupCount().by("title")

```
SELECT
b
MATCH
(a) -[:created]-> (b)
WHERE
a.name = "gremlin"
GROUP

BY
c.title
```



Gremlin

```
g.V()
    .as("a")
    .has("name", "gremlin")
    .out("created")
    .in("created")
    where(neq("a"))
    groupCount().by("title")
```

```
SELECT
   c
MATCH
   (a) -[:created]-> (b),
   (b) <-[:created]- (c)
WHERE
   a.name = "gremlin"
GROUP</pre>
BY
```

Gremlin

```
g.V()
    .as("a")
    .has("name", "gremlin")
    .out("created")
    .in("created")
    .where(neq("a"))
    groupCount().by("title")
```

```
SELECT
   c
MATCH
   (a) -[:created]-> (b),
   (b) <-[:created]- (c)
WHERE
   c != a
AND
   a.name = "gremlin"
GROUP BY
   c.title</pre>
```



Gremlin

```
g.V()
    .as("a")
    .has("name", "gremlin")
    .out("created")
    .in("created")
    .where(neq("a"))
    .groupCount().by("title")
```

```
SELECT
   c.title, COUNT(*)

MATCH
   (a) -[:created]-> (b),
   (b) <-[:created]- (c)

WHERE
   c != a

AND
   a.name = "gremlin"

GROUP BY
   c.title</pre>
```



```
martin@martin-arch:gremlin2pgx(pgql)$ /usr/lib/jvm/java-11-openjdk/bin/java -cp $(cat gremlin2pgx.shell/build/class
path.txt) oracle
.gm.lang.gremlin2pgx.Shell
| Welcome to JShell -- Version 11.0.4
| For an introduction type: /help intro
```

jshell> ■

```
martin@martin-arch:gremlin2pgx(pgql)$ /usr/lib/jvm/java-11-openjdk/bin/java -cp $(cat gremlin2pgx.shell/build/class
path.txt) oracle
.gm.lang.gremlin2pgx.Shell
| Welcome to JShell -- Version 11.0.4
| For an introduction type: /help intro
```

jshell> g.V().out()■

```
martin@martin-arch:gremlin2pgx(pgql)$ /usr/lib/jvm/java-11-openjdk/bin/java -cp $(cat gremlin2pgx.shell/build/class
path.txt) oracle
.gm.lang.gremlin2pgx.Shell
| Welcome to JShell -- Version 11.0.4
| For an introduction type: /help intro

jshell> g.V().out()
$11 ==> [ GraphStep(vertex, []), VertexStep(OUT, vertex) ]
```

jshell>

```
martin@martin-arch:gremlin2pgx(pgql)$ /usr/lib/jvm/java-11-openjdk/bin/java -cp $(cat gremlin2pgx.shell/build/class
path.txt) oracle
.gm.lang.gremlin2pgx.Shell
| Welcome to JShell -- Version 11.0.4
| For an introduction type: /help intro

jshell> g.V().out()
$11 ==> [ GraphStep(vertex, []), VertexStep(OUT, vertex) ]
```

jshell> t(g.V().out())■

```
martin@martin-arch:gremlin2pgx(pgql)$ /usr/lib/jvm/java-11-openjdk/bin/java -cp $(cat gremlin2pgx.shell/build/class
path.txt) oracle
.gm.lang.gremlin2pgx.Shell
| Welcome to JShell -- Version 11.0.4
| For an introduction type: /help intro

jshell> g.V().out()
$11 ==> [ GraphStep(vertex, []), VertexStep(OUT, vertex) ]

jshell> t(g.V().out())
$12 ==> "SELECT v1 MATCH (v0) -> (v1)"

jshell> ■
```

```
martin@martin-arch:gremlin2pgx(pgql)$ /usr/lib/jvm/java-11-openjdk/bin/java -cp $(cat gremlin2pgx.shell/build/class
path.txt) oracle
.gm.lang.gremlin2pgx.Shell
  Welcome to JShell -- Version 11.0.4
  For an introduction type: /help intro
jshell> g.V().out()
$11 ==> [ GraphStep(vertex, []), VertexStep(OUT, vertex) ]
jshell> t(g.V().out())
$12 ==> "SELECT v1 MATCH (v0) -> (v1)"
jshell> (g.V()
  ...> .hasLabel("Customer").as("c")
  ...> .values("name").as("name")
        .select("c")
  . . . >
          .out("Bought").as("b")
  . . . >
          .values("discount").as("disc")
  ...>
         .select("b")
  . . . >
         .in()
  . . .>
         .hasLabel("Product")
  . . . >
```

...>)

```
martin@martin-arch:gremlin2pgx(pgql)$ /usr/lib/jvm/java-11-openjdk/bin/java -cp $(cat gremlin2pgx.shell/build/class
path.txt) oracle
.gm.lang.gremlin2pgx.Shell
Welcome to JShell -- Version 11.0.4
  For an introduction type: /help intro
jshell> g.V().out()
$11 ==> [ GraphStep(vertex, []), VertexStep(OUT, vertex) ]
jshell> t(g.V().out())
$12 ==> "SELECT v1 MATCH (v0) -> (v1)"
jshell> (g.V()
  ...> .hasLabel("Customer").as("c")
  ...> .values("name").as("name")
  ...> .select("c")
         .out("Bought").as("b")
  . . . >
         .values("discount").as("disc")
  . . . >
  . . . >
        .select("b")
        .in()
  . . . >
         .hasLabel("Product")
  . . . >
  ...>)
$13 ==> [ GraphStep(vertex, []), HasStep([~label.eq("Customer")]), AsStep(c), PropertiesStep([name], value), AsStep
(name), SelectOneStep(last, c), VertexStep(OUT, [Bought], vertex), AsStep(b), PropertiesStep([discount], value), As
Step(disc), SelectOneStep(last, b), VertexStep(IN, vertex), HasStep([~label.eg("Product")]) ]
jshell>
```

```
Welcome to JShell -- Version 11.0.4
  For an introduction type: /help intro
jshell> g.V().out()
$11 ==> [ GraphStep(vertex, []), VertexStep(OUT, vertex) ]
jshell> t(g.V().out())
$12 ==> "SELECT v1 MATCH (v0) -> (v1)"
jshell> (g.V()
  ...> .hasLabel("Customer").as("c")
  ...> .values("name").as("name")
  ...> .select("c")
        .out("Bought").as("b")
  . . . >
  ...>
         .values("discount").as("disc")
        .select("b")
  . . .>
  . . .>
         .in()
         .hasLabel("Product")
  . . . >
  ...>)
$13 ==> [ GraphStep(vertex, []), HasStep([~label.eq("Customer")]), AsStep(c), PropertiesStep([name], value), AsStep
(name), SelectOneStep(last, c), VertexStep(OUT, [Bought], vertex), AsStep(b), PropertiesStep([discount], value), As
Step(disc), SelectOneStep(last, b), VertexStep(IN, vertex), HasStep([~label.eq("Product")]) ]
jshell> t(g.V()
          .hasLabel("Customer").as("c")
          .values("name").as("name")
  . . .>
  . . .>
         .select("c")
  . . .>
          .out("Bought").as("b")
          .values("discount").as("disc")
  . . .>
          .select("b")
  . . .>
  . . .>
          .in()
           .hasLabel("Product")
  . . . >
```

...>)

```
$11 ==> [ GraphStep(vertex, []), VertexStep(OUT, vertex) ]
jshell> t(g.V().out())
$12 ==> "SELECT v1 MATCH (v0) -> (v1)"
jshell>(g.V()
  ...> .hasLabel("Customer").as("c")
         .values("name").as("name")
  ...> .select("c")
        .out("Bought").as("b")
  . . . >
  ...>
         .values("discount").as("disc")
         .select("b")
  . . .>
  . . .>
        .in()
         .hasLabel("Product")
  ...>
  ...>)
13 ==  [ GraphStep(vertex, []), HasStep([~label.eq("Customer")]), AsStep(c), PropertiesStep([name], value), AsStep
(name), SelectOneStep(last, c), VertexStep(OUT, [Bought], vertex), AsStep(b), PropertiesStep([discount], value), As
Step(disc), SelectOneStep(last, b), VertexStep(IN, vertex), HasStep([~label.eq("Product")]) ]
jshell> t(g.V()
  ...> .hasLabel("Customer").as("c")
          .values("name").as("name")
  . . .>
  ...>
         .select("c")
         .out("Bought").as("b")
  . . .>
          .values("discount").as("disc")
  . . .>
  . . .>
          .select("b")
  . . .>
          .in()
          .hasLabel("Product")
  ...>
  ...> )
$14 ==> "SELECT v2 MATCH (v0:Customer) -[:Bought]-> (v1), (v1) <- (v2:Product)"
jshell>
```

jshell> g.V().out()

Supported Gremlin Features

map

- Constant
- Count
- Id
- Label
- Match
- Math
- Max
- Mean
- Min
- Order
- Project

- Select
- Sum

flatMap

Vertex

filter

- And
- Dedup
- Has
- Is
- Limit
- Not

- Or
- Range

- Skip
- Where

sideEffect

- Group
- GroupCount
- ...

Unsupported Gremlin Features

- Multiple input graphs
 - Rarely seen in real-life queries
- Graph mutation (insert / delete)
 - Will be supported in future PGQL version
- Control flow / algorithmic computation
 - Oracle's Spatial & Graph offers the PGX Algorithm API
 - Intuitive Java API for implementing graph algorithms
 - Offer specialized features for each use-case
 - PGQL for queries
 - PGX Algorithm API for algorithms

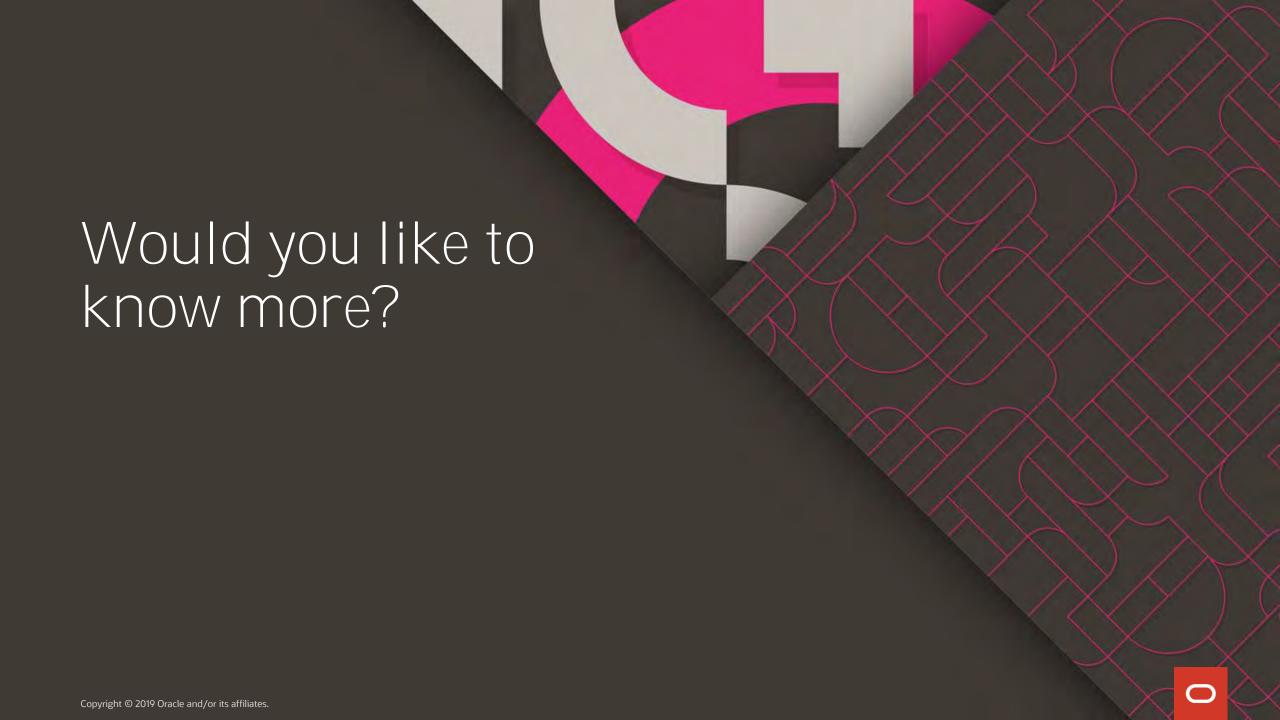


Summary

- PGQL is Oracle's property graph query language
 - Powerful pattern matching query language
 - Has integration with graph algorithms
 - Part of Oracle [Big Data] Spatial and Graph
- Gremlin is a popular open source graph traversal language
- Research area: We can auto-transform Gremlin to PGOL*
 - Easier transition for users
 - Utilize benefits of PGQL



^{*}Currently not included in shipping product – Oracle Safe Harbor terms apply



Graph at OOW and Code One 2019

<u>Sessions</u>

Wednesday, Sept. 18			
10:00 am - 10:45 am	Graph Databases and Analytics: How To Use Them [TRN4755]	Moscone South - Room 152C	
10:00 am - 10:45 am	Setting Up Modern Anti-Money-Laundering Solutions to Service Wholesale [CON6223]	Moscone West - Room 3004	
11:15 am - 12:00 pm	Demystifying Graph Analytics for the Nonexpert [CON5503]	Moscone South - Room 156B	
1:30 pm - 2:15 pm	Traversing and Querying Graphs with PGQL and Gremlin with Oracle Spatial and Graph [DEV4084]	Moscone South - Room 202	

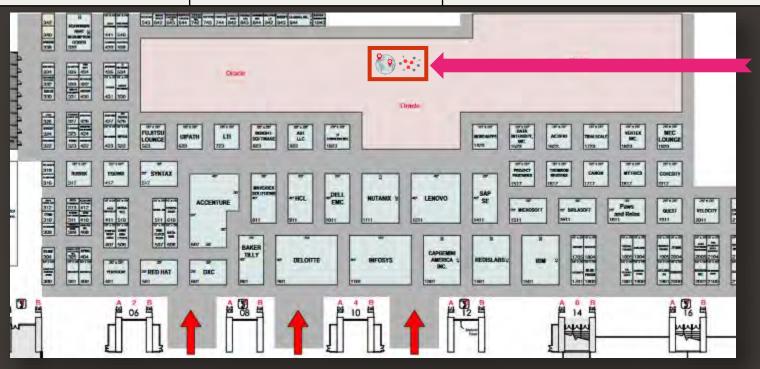
Meet the Experts At the Code One Groundbreakers Hub, Moscone South Level 1

Wednesday, Sept. 18		
1:30 pm - 2:20 pm	Graph Database and Analysis	Lounge C, Code One Groundbreakers Hub, Moscone South level 1
2:30 pm - 3:20 pm	Graph Cloud Service: Automating Graph Analysis	



Graph at OOW and Code One 2019

Date/TimeTitleLocationMonday 10:00 am – 4:00 pm
Tuesday 10:30 am – 5:30 pm
Wednesday 10:00 am – 4:30 pmSpatial and Graph:
Database, Analytics and
CloudMoscone South Exhibit Hall ('The Exchange')
Oracle Demogrounds > Data Management area
> Kiosk # ODB-017



Spatial & Graph





Sungpack Hong

Research Director

Oracle Labs

Martin Sevenich

Principal Researcher

Oracle Labs

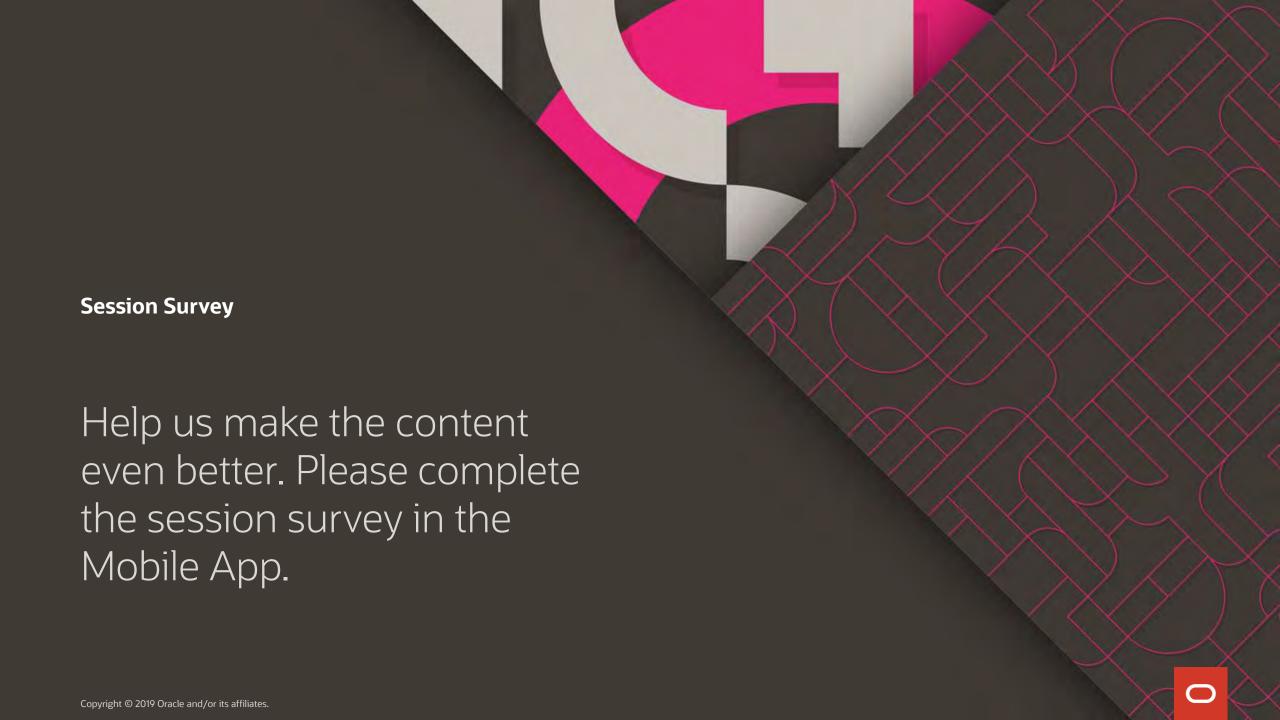
Oskar van Rest

Principal Member of Technical Staff Oracle Spatial and Graph

Martijn Dwars

Senior Researcher

Oracle Labs



Safe Harbor

The preceding is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, timing, and pricing of any features or functionality described for Oracle's products may change and remains at the sole discretion of Oracle Corporation.

Statements in this presentation relating to Oracle's future plans, expectations, beliefs, intentions and prospects are "forward-looking statements" and are subject to material risks and uncertainties. A detailed discussion of these factors and other risks that affect our business is contained in Oracle's Securities and Exchange Commission (SEC) filings, including our most recent reports on Form 10-K and Form 10-Q under the heading "Risk Factors." These filings are available on the SEC's website or on Oracle's website at http://www.oracle.com/investor. All information in this presentation is current as of September 2019 and Oracle undertakes no duty to update any statement in light of new information or future events.



Break New Ground

San Francisco September 16–19, 2019

