



Daniël ten Wolde

#### The State of DuckPGQ

Peter Boncz (on sabbatical @ <br/>
MotherDuck)<br/>
CWI Database Architectures group

## Graph data management



#### Graph data management





#### Graph data management



#### graph exploration





SELECT count(\*)
FROM person
WHERE name LIKE 'E%'

#### relational operators

# Storing graphs in SQL

```
CREATE TABLE city (
                                                                    :person
                                                                   name: Bob
  id bigint PRIMARY KEY,
  name varchar
                                                          :person
                                                                             :person
);
                                                         name: Chloe
                                                                            name: Jack
CREATE TABLE person (
                                                       follows
                                                                    :person
                                                                                          :city
  id bigint PRIMARY KEY,
                                                                  name: Emily
                                                                                      name: Utrecht
  name varchar,
                                                                              livesIn
  livesIn bigint,
  CONSTRAINT c FOREIGN KEY (livesIn) REFERENCES city (id)
);
CREATE TABLE follows (
  p1id bigint,
  p2id bigint,
  CONSTRAINT p1 FOREIGN KEY (p1id) REFERENCES person (id),
  CONSTRAINT p2 FOREIGN KEY (p2id) REFERENCES person (id)
);
```

#### SQL:1999 query

```
WITH RECURSIVE paths(startNode, endNode, path) AS (
  SELECT plid AS startNode, p2id AS endNode, ARRAY[p1id, p2id] AS path
    FROM follows JOIN person p1 ON p1.id = follows.p1id WHERE p1.name = 'Bob'
  UNION ALL (
    WITH paths AS (TABLE paths)
       SELECT paths.startNode AS startNode, p2id AS endNode, array_append(path, p2id) AS path
       FROM paths JOIN follows ON paths.endNode = follows.p1id
       WHERE NOT EXISTS (SELECT true FROM paths previous_paths
                         JOIN person p2 ON p2.id = follows.p2id
                          WHERE p2.name = 'Bob' OR follows.p2id = previous_paths.endNode)))
SELECT count(p2.id) AS cp2
FROM person p1
JOIN paths ON paths.startNode = p1.id
JOIN person p2 ON p2.id = paths.endNode
JOIN city ON city.id = p2.livesIn AND city.name = 'Utrecht'
```

#### SQL:1999 query

```
WITH RECURSIVE paths(startNode, endNode, path) AS (
  SELECT plid AS startNode, p2id AS endNode, ARRAY[p1id, p2id] AS path
    FROM follows JOIN person p1 ON p1.id = follows.p1id WHERE p1.name = 'Bob'
  UNION ALL (
    WITH paths AS (TABLE paths)
       SELECT paths.startNode AS startNode, p2id AS endNode, array_append(path, p2id) AS path
       FROM paths JOIN follows ON paths.endNode = follows.p1id
       WHERE NOT EXISTS (SELECT true FROM paths previous_paths
                         JOIN person p2 ON p2.id = follows.p2id
                          WHERE p2.name = 'Bob' OR follows.p2id = previous_paths.endNode)))
SELECT count(p2.id) AS cp2
FROM person p1
JOIN paths ON paths.startNode = p1.id
JOIN person p2 ON p2.id = paths.endNode
JOIN city ON city.id = p2.livesIn AND city.name = 'Utrecht'
```

#### SQL:1999 query

```
WITH RECURSIVE paths(startNode, endNode, path) AS (
  SELECT p1id AS startNode, p2id AS endNode, ARRAY[p1id, p2id] AS path
    FROM follows JOIN person p1 ON p1.id = follows.p1id WHERE p1.name = 'Bob'
  UNION ALL (
    WITH paths AS (TABLE paths)
       SELECT paths.startNode AS startNode, p2id AS endNode, array_append(path, p2id) AS path
       FROM paths JOIN follows ON paths.endNode = follows.p1id
       WHERE NOT EXISTS (SELECT true FROM paths previous_paths
                         JOIN person p2 ON p2.id = follows.p2id
                          WHERE p2.name = 'Bob' OR follows.p2id = previous_paths.endNode)))
SELECT count(p2.id) AS cp2
FROM person p1
JOIN paths ON paths.startNode = p1.id
JOIN person p2 ON p2.id = paths.endNode
JOIN city ON city.id = p2.livesIn AND city.name = 'Utrecht'
```



## Graph query languages

Most popular graph system **,∩eO4j** 

Cypher



SPARQL/Gremlin



Gremlin

Oracle Labs PGX

**PGQL** 



GSQL



nGQL

#### SQL/PGQ (Property Graph Queries)



#### SQL/PGQ

- Extension in the SQL:2023 standard, released in June 2023
- "Property Graphs" defined over existing tables
- Read-only operations for graph queries return "Graph Tables"
  - Path-finding
  - Pattern matching

#### Tabular schema

```
CREATE TABLE city (
  id bigint PRIMARY KEY,
  name varchar
);
CREATE TABLE person (
  id bigint PRIMARY KEY,
  name varchar,
  livesIn bigint,
  CONSTRAINT c FOREIGN KEY ...
);
CREATE TABLE follows (
  p1id bigint,
  p2id bigint,
  CONSTRAINT p1 FOREIGN KEY ...
 CONSTRAINT p2 FOREIGN KEY ...
);
```

# SQL/PGQ graph tables

```
CREATE PROPERTY GRAPH socialNetwork
VERTEX TABLES (
city,
person
)
EDGE TABLES (
livesIn SOURCE person DESTINATION city,
follows SOURCE person DESTINATION person
);
```



Prompt: Count the number of people Bob (in)directly follows who live in the city Utrecht

```
SELECT count(id)
FROM
GRAPH_TABLE (socialNetwork,
```

MATCH (p1:person WHERE p1.name='Bob')-[:follows]->\*(p2:person)
 -[:livesIn]->(c:city WHERE c.name='Utrecht')

**COLUMNS** (p2.id)

plain

SQL

```
WITH RECURSIVE paths(startNode, endNode, path) AS (
  SELECT p1id AS startNode, p2id AS endNode, ARRAY[p1id, p2id] AS path
    FROM follows JOIN person p1 ON p1.id = follows.p1id WHERE p1.name = 'Bob'
  UNION ALL (
   WITH paths AS (TABLE paths)
      SELECT paths.startNode AS startNode, p2id AS endNode, array_append(path, p2id) AS path
      FROM paths JOIN follows ON paths.endNode = follows.p1id
      WHERE NOT EXISTS (SELECT true FROM paths previous_paths
                       JOIN person p2 ON p2.id = follows.p2id
                        WHERE p2.name = 'Bob' OR follows.p2id = previous_paths.endNode)))
SELECT count(p2.id) AS cp2
FROM person p1
JOIN paths ON paths.startNode = p1.id
JOIN person p2 ON p2.id = paths.endNode
JOIN city ON city.id = p2.livesIn AND city.name = 'Utrecht'
```

#### The SQL/PGQ query is 4× more concise

















#### Implementation of SQL/PGQ in DuckDB





# **DuckDB: in-process analytics**

- Created by Hannes Mühleisen and Mark Raasveldt
- Idea: analytical SQL system as a linkable library



- From research on data systems support for data science:
  - why don't data scientists use database systems? Ο
    - $\Rightarrow$  make database technology better suited for data science
- Active discord, blog, starting events, traction:
  - 18K github stars, >2M downloads/month (5x increase YoY) Ο
  - DuckDB Labs spin-off (+MotherDuck) Ο





#### **DuckDB Extension Framework**



#### **DuckDB Extension Framework**



In the past decade, property graph databases have emerged as a

views even relational tables and (2) to formulate much

## DuckPGQ: extension module for DuckDB

 Duck parser extension rewrites queries + UDFs

CWI

Centrum Wiskunde & Informatica



#### **DuckPGQ: extension module for DuckDB**

- Duck **parser extension** rewrites queries + UDFs

CWI

Centrum Wiskunde & Informatica

- CSR creation on-the-fly, exploiting ROWIDs to get dense node-IDs quickly
- Scalar UDFs: multi-core **parallelism** out of the box





#### **Compressed Sparse Row (CSR) data structure**

- On-the-fly creation
- Two scalar UDFs
  - Initialize vertex array
  - Initialize edge array
- Index in the **vertex array** corresponds to the row identifier of the vertex
- Vertex array contains offsets for the edge arrays



3

4

5

10

11

12

#### **DuckPGQ: extension module for DuckDB**

- Duck **parser extension** rewrites queries + UDFs

CWI

Centrum Wiskunde & Informatica

- CSR creation on-the-fly, exploiting ROWIDs to get dense node-IDs quickly
- Scalar UDFs: multi-core **parallelism** out of the box



#### **DuckPGQ: extension module for DuckDB**

- Duck **parser extension** rewrites queries + UDFs

CWI

Centrum Wiskunde & Informatica

- CSR creation **on-the-fly**, exploiting ROWIDs to get dense node-IDs quickly
- Scalar UDFs: multi-core **parallelism** out of the box
- SIMD-efficient Multi-Source BFS (and pathfinding algos)





# **MS-BFS** Algorithms

Multiple-Source (MS) graph algorithms

Idea:

- Do many (512 or more) searches at-a-time
   "vectorized"
- Keep state for many (512 or more) searches
- Store state in **SIMD registers**
- Algorithm: sequential access Vertex & Edge arrays
  - Random access in Vertex for destination state
  - but it is shared for 512 searches



Figure 3: An example showing the steps of MS-BFS when using bit operations. Each row represents the bit field for a vertex, and each column corresponds to one BFS. The symbol X indicates that the value of the bit is 1.

#### **PVLDB 8 (2014)**

The More the Merrier: Efficient Multi-Source Graph Traversal						
Manuel Then*	Moritz Kaufmann*	Fernando Chirigati <sup>†</sup>	Tuan-Anh Hoang-Vu <sup>†</sup>			
then@in.tum.de	kaufmanm@in.tum.de	fchirigati@nyu.edu	tuananh@nyu.edu			
Kien Pham <sup>†</sup>	Alfons Kemper*	Thomas Neumann*	Huy T. Vo <sup>†</sup>			
en.pham@nyu.edu	kemper@in.tum.de	neumann@in.tum.de	huy.vo@nyu.edu			
* Technische Universität München		<sup>†</sup> New York University				

ABSTRACT

Graph analytics on social networks, Web data, and communication networks has been widely used in a plethora of applications. Many graph analytics algorithms are based on breadth-first search (BFS) graph traversal, which is not only ime-consuming for large datasets but also involves much redundant computation when executed multiple times from different start vertices. In this paper, we propose *Multi-Source BFS* (MS-BFS), an algorithm that is designed to run multiple concurrent BFSs over the same graph on a sinele CPU core while scaling un as the number of cores have influence on others and, as a consequence, are of great importance to spread information, e.g., for marketing purposes [20].

In a wide range of graph analytics algorithms, including shortest path computation [13], graph centrality calculation [9, 27], and k-hop neighborhood detection [12], breadthfirst search (BFS)-based graph traversal is an elementary building block used to systematically traverse a graph, i.e., to visit all reachable vertices and edges of the graph from a given start vertex. Because of the volume and nature of the data, BFS is a computationally expensive operation. lead-



#### **DuckPGQ Extensions: weighted path**

MATCH ANY SHORTEST PATHS p=(a:Person)-[e:know]->+(b:Person)
COLUMNS (ELEMENT\_ID(a) aid, a.name src,

ELEMENT\_ID(b) bid, b.name dst, COST(p), p))

aid	src	bid	dst	COST	p
int64	varchar	int64	varchar	int32	int64[]
0 0 2 3 3 0 3	Ana Ana Ed Jo Jo Ana Jo	1 3 1 0 2 2 1	Bo Jo Bo Ana Ed Ed Bo	1 1 1 2 2	$\begin{bmatrix} 0, & 0, & 1 \\ [0, & 1, & 3 ] \\ [2, & 2, & 1 ] \\ [3, & 3, & 0 ] \\ [3, & 4, & 2 ] \\ [0, & 1, & 3, & 4, & 2 ] \\ [3, & 4, & 2, & 2, & 1 ] \end{bmatrix}$

Note: the green numbers are ELEMENT IDs of edges

^



#### **DuckPGQ Extensions: label masks**

```
CREATE PROPERTY GRAPH pg
VERTEX TABLES (
 College PROPERTIES (id, college) LABEL College,
 Person PROPERTIES (id, name, birthDate) LABEL Person
                                          IN msk (Student, TA))
EDGE TABLES (
 know SOURCE KEY(src) REFERENCES Person(id)
       DESTINATION KEY(dst) REFERENCES Person(id)
       PROPERTIES (createDate, msgCount) LABEL know,
 enrol SOURCE
                  KEY(studentID) REFERENCES Person(id)
       DESTINATION KEY(collegeID) REFERENCES College(id)
       PROPERTIES (classYear) LABEL studiesAt );
```

## **Ongoing work: GNN integration**

- Analyze PGQ property graphs in DGL and Pytorch Geometric
- Export DGL and PyTorch Geometric graphs to PGQ property graphs

```
import dgl
1
    import torch
2
    import duckdb
3
4
    # Setup DuckPGQ and collect necessary data
5
    con = duckdb.connect()
6
    csrv, csre, node_features, edge_features = ...
7
    # Initialize a graph object
8
    g = dgl.graph(('csr', (csrv, csre, [])))
9
    # Set the node features, reshaping is necessary
10
    for feature name, feature in node features:
11
        g.ndata[feature_name] = feature.reshape((feature.shape[0], 1))
12
    # Set the edge features, reshaping is necessary
13
    for feature_name, feature in edge_features:
14
        g.edata[feature_name] = feature.reshape((feature.shape[0], 1))
15
```

#### Ongoing work in DuckPGQ



#### Ongoing work: Parallel Pathfinding

- Scalar function **find\_path(src,dst)** has limitations
  - *Morsel-driven* parallelism on [src,dst] table
  - morsel=120K tuples that's a lot of searches!



#### **Ongoing work: Parallel Pathfinding**

- Scalar function **find\_path(src,dst)** has limitations
  - *Morsel-driven* parallelism on [src,dst] table
  - morsel=120K tuples that's a lot of searches!
- New project: DuckDB pathfinding operator
  - Every BFS frontier advance is a DuckDB event
  - Threads are **scheduled** to work on vertex-ranges
  - **Source** starts with [src,dst] materialization
    - deduplication of src, dst, [src,dst])
  - Sink re-creates the found paths
    - proper order and duplicity

#### **Fine-grained Parallelism**



- DuckDB for joins uses a bucket-chained hash-table
  - Mixes hash-conflicts with duplicates in a chain
- We changed the hash-table to only have chains for duplicates
  - Linear hashing on the buckets
  - Its' faster on joins!

- DuckDB for joins uses a bucket-chained hash-table
  - Mixes hash-conflicts with duplicates in a chain
- We changed the hash-table to only have chains for duplicates
  - Linear hashing on the buckets
  - Its' faster on joins!

#### **CIDR 2022**

#### The 3D Hash Join: Building On Non-Unique Join Attributes

Daniel Flachs flachs@uni-mannheim.de University of Mannheim Mannheim, Germany

#### ABSTRACT

One of the most prominent ways to evaluate an equi-join is based on hashing. We consider the problem of non-unique join attributes on the build side. In conventional hash tables where collisions are resolved by chaining, duplicates inevitably lead to long collision chains. This causes a high number of expensive main memory

Magnus MüllerGuido Moerkottemagnus@uni-mannheim.demoerkotte@uni-mannheim.deUniversity of MannheimUniversity of MannheimMannheim, GermanyMannheim, Germany

when traversing the collision chains, resulting in high processing costs for probing.

A related problem occurs if the uniqueness of the join attributes in the build relation is not known at query compilation time. This happens if the *known functional dependencies* specified in the SQL standard do not allow to derive uniqueness.

- Joins can now return a hit-list
  - A **hit-list** points to a hash-table chain  $\Rightarrow$  **factorized n:m join**
- Exploiting these:
  - Factorized Aggregation: Embedding (sub-)aggregates in hit-lists
  - Factorized Joins: Embedding hit-lists in hit-lists ("graph construction")
  - Worst-Case Optimal Joins (WCOJ): Cyclical joins using hit-list intersection

- Joins can now return a **hit-list** 
  - A hit-list points to a hash-table chain ⇒ factorized n:m join
- Exploiting these:
  - Factorized Aggregation: Embedding (sub-)aggregates in hit-list
  - Factorized Joins: Embedding hit-lists in hit-lists ("graph construction")
  - Worst-Case Optimal Joins (WCOJ): Cyclical joins using hit-list intersection

**Adaptive Query** 

execution

# Conclusion

- SQL/PGQ (skipped)
- DuckPGQ last year
  - Parser support + scalar UDF MS-BFS pathfinding
  - Extensions: weighted shortest path-ding, flexible labels
- DuckPGQ ongoing work
  - GNN library integrations
  - MS-BFS operator (better parallelism)
  - Factorized query processing



- duckdb -unsigned
- v1.0.0 1f98600c2c
- Enter ".help" for usage hints.
- D set custom\_extension\_repository = 'http://duckpgq.s3.eu-north-1.amazonaws.com';
- D force install 'duckpgq';
- D load 'duckpag':



Daniël ten Wolde

DuckPGQ on GitHub