A SQL OLAP benchmark based on the LDBC Social Network Benchmark

Nuno Faria

INESCTEC GDC

Graph Data Council Twenty-First TUC Meeting - Boston, Sep 2026

Motivation

- Large sets of data are commonly stored as files (e.g., Parquet) in data lakes, like cloud-based object stores.
- It might not be feasible to pre-load large files to a database to perform analytical workloads.
- Analytical engines that execute directly over these data files (e.g., employing projection and filter pushdown) commonly offer a SQL interface.
- It would be interesting to explore how these systems handle graph-based analytical workloads.

Social Network Benchmark

- Models the data and interactions in a social network. Main entities include persons, forums, and messages.
- Supports two main types of workloads:
  - Interactive
  - **Business Intelligence (BI)**
- BI workload contains 20 analytical queries, each modeling different choke points (e.g., worst-case optimal joins).

Social Network Benchmark schema

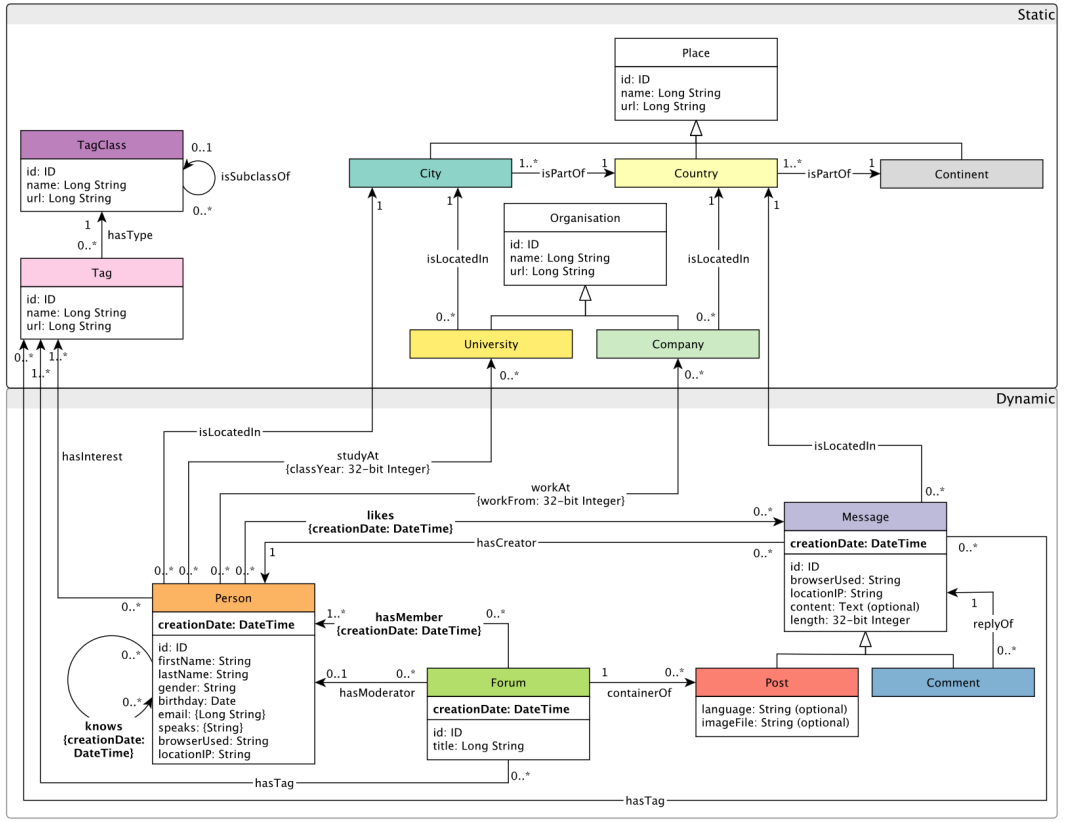

Source: The LDBC Social Network Benchmark (version 2.2.4). https://arxiv.org/pdf/2001.02299

Engine Requirements

- SQL analytical engine.
- Can execute directly over Parquet files.
- Ideally supports recursive queries.
  - In the SNB BI benchmark, three queries need recursive support (Q15, Q19, Q20).
- Can be locally deployed.

Recursive Queries

- Allow SQL to execute queries over graph-like structures.
- Implemented with recursive Common Table Expressions (CTEs).
- Turn SQL into a Turing complete language (https://github.com/nuno-faria/tetris-sql)

Generator Query

```
WITH RECURSIVE t AS
(
  SELECT 1 AS i

  UNION ALL

  SELECT i + 1
  FROM t
  WHERE i < 5
)
SELECT *
FROM t;
```

i
1
2
3
4
5

Recursive Queries

- **Example:** Bill of Materials

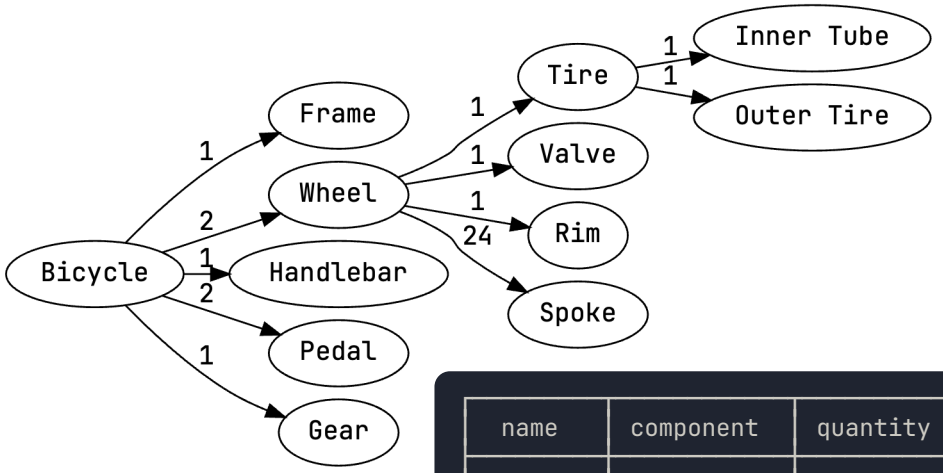

name	component	quantity
Bicycle	Frame	1
Bicycle	Wheel	2
Bicycle	Handlebar	1
Bicycle	Pedal	2
Bicycle	Gear	1
Wheel	Tire	1
Wheel	Valve	1
Wheel	Rim	1
Wheel	Spoke	24
Tire	Inner Tube	1
Tire	Outer Tire	1

Cypher

```
MATCH (b:Part {name: 'Bicycle'})
  -[rels:CONTAINS]→+(m:Part)
RETURN m.name AS component,
  reduce(q = 1, r IN rels | q * r.quantity) AS quantity
ORDER BY component;
```

SQL

```
WITH RECURSIVE m AS (
  SELECT component, quantity
  FROM part
  WHERE name = 'Bicycle'
  UNION ALL
  SELECT p.component, m.quantity * p.quantity
  FROM m
  JOIN part p ON p.name = m.component
)
SELECT *
FROM m
ORDER BY component;
```

Survey of Recursive CTE Support

Engine	Supports recursive CTEs?
cedardb	Yes
clickhouse	UNION not supported / Q19 times out
cloudberry	SNB recursive queries crash
databend	UNION not supported / Q19 OOM
datafusion	Cannot be referenced multiple times
doris	Cannot be referenced multiple times
dremio	No
duckdb	Yes
firebolt	No
hive	No
impala	No
oceanbase	Cannot be referenced multiple times
parquet_fdw	Yes
polars	No
presto	No
questdb	No
spark	UNION not supported / OOM / Other limitations
starrocks	Cannot be referenced multiple times / Too slow
trino	Cannot be referenced multiple times / Stack overflow

- The recursive queries in SNB BI (Q15, Q19, Q20) reference recursive CTEs multiple times. Q15 and Q20 also require UNION.

Benchmark Description

- Each system registers the same Parquet files (no data pre-loading).
- 20 SNB BI queries are executed (average of multiple runs).
- Additional queries are created and evaluated to test specific behaviors.

Preliminary Results

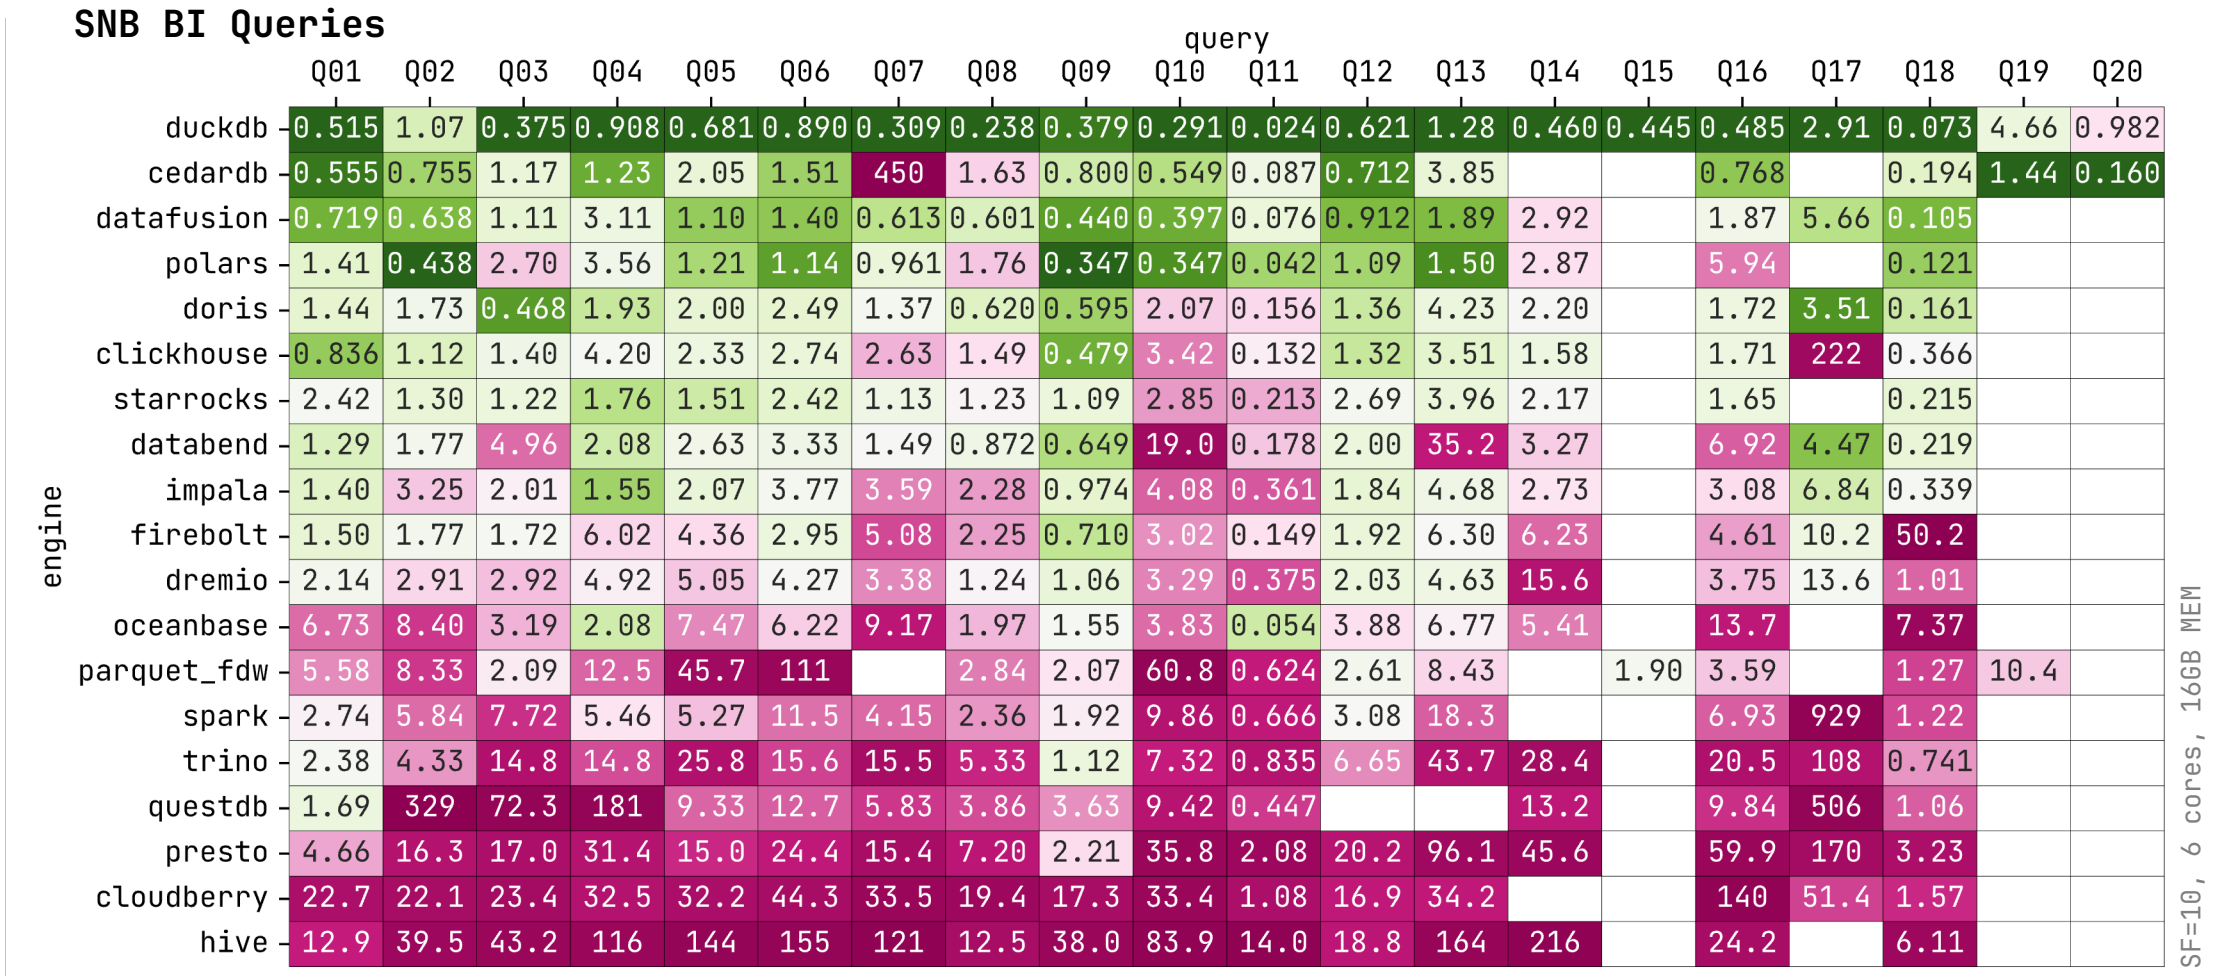

Preliminary Results

- Preliminary conclusions:
  - Q17 is overall the slowest query to execute. Many systems select a sub-optimal join order.

Excerpt of Q17 in ClickHouse

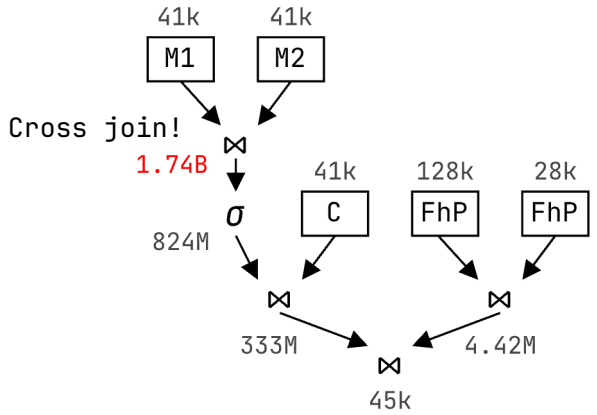

Excerpt of Q17 in DuckDB

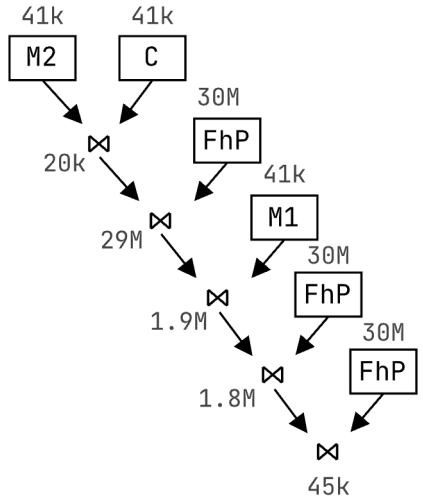

- Bad join ordering is also why Q07 in CedarDB is very expensive.
- DuckDB is the only engine able to execute the full workload.
- There is a large performance disparity between the different systems.

Preliminary Results

- Selected additional queries:

Sequence generator **with a recursive CTE (1M iters)**

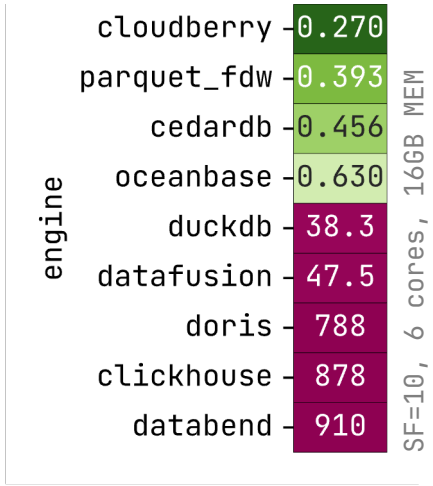

Connected components **of a sample of Person_knows_Person relationships**

10k relationships

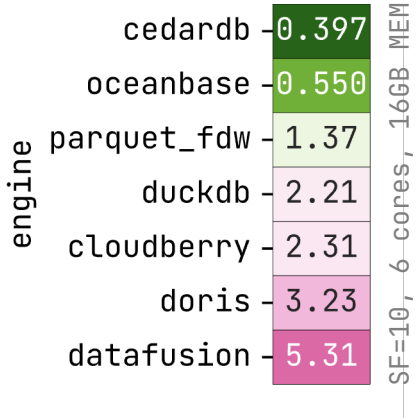

32k relationships

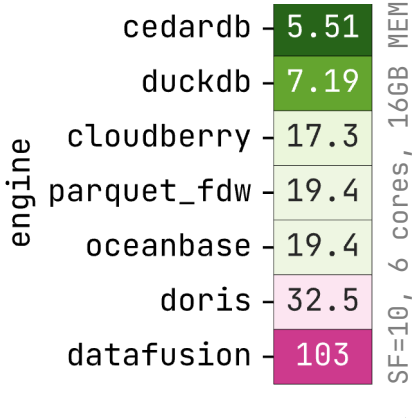

Single-Source Shortest Paths **(sample of 110k relationships, positive weights)**

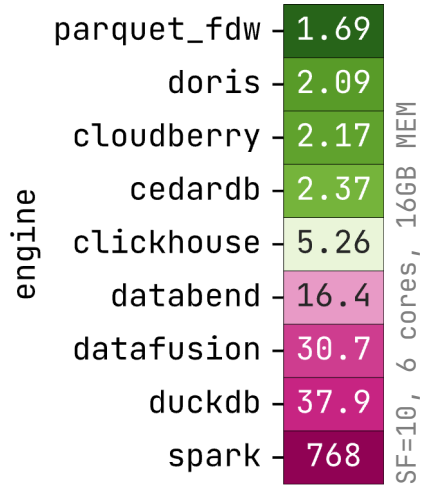

Preliminary Results

- The “single-source shortest paths” query can be improved with the “USING KEY” feature (currently only available in DuckDB).
- Allows recursive queries to access the result (“RECURRING.t”) and update its contents.

Bamberg, Björn, Denis Hirn, and Torsten Grust. "How DuckDB is USING KEY to Unlock Recursive Query Performance." Companion of the 2025 International Conference on Management of Data. 2025. https://dl.acm.org/doi/pdf/10.1145/3722212.3725107

```
WITH RECURSIVE t USING KEY (k) (
   ...
)
```

Initial graph

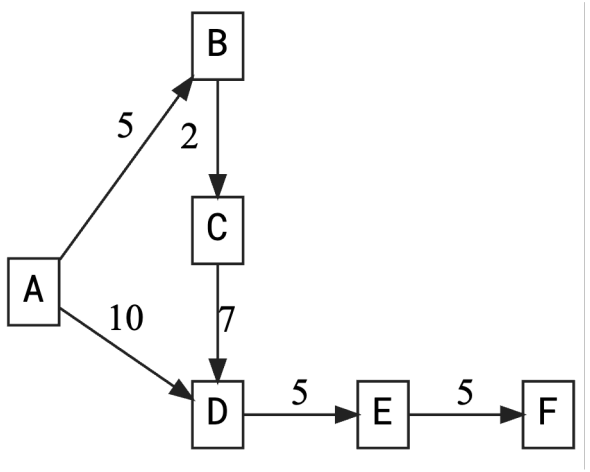

Regular recursive CTE

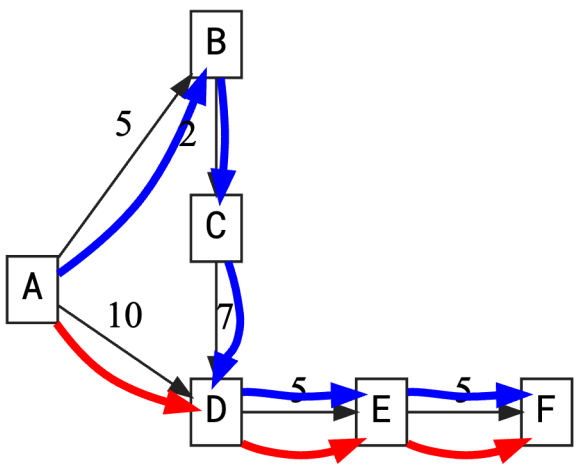

Recursive CTE + USING KEY

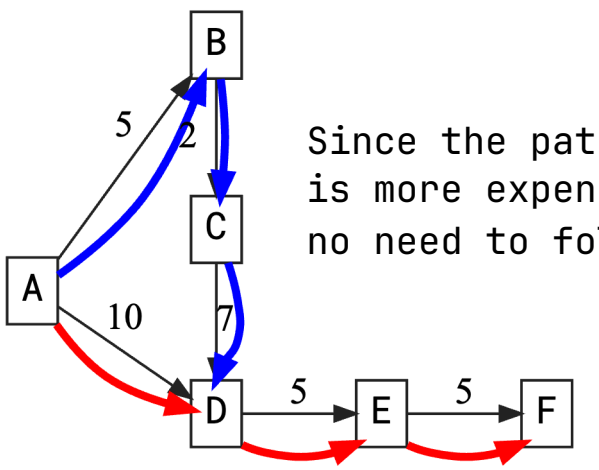

Since the path to D from C is more expensive, there is no need to follow it.

Preliminary Results

- Single-source shortest paths without (left) and with USING KEY (right) in DuckDB (110k relationships):

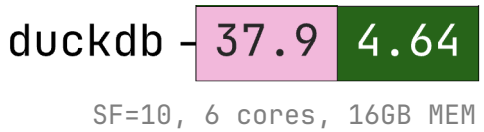

- DuckDB vs DuckDB USING KEY vs parquet_fdw:

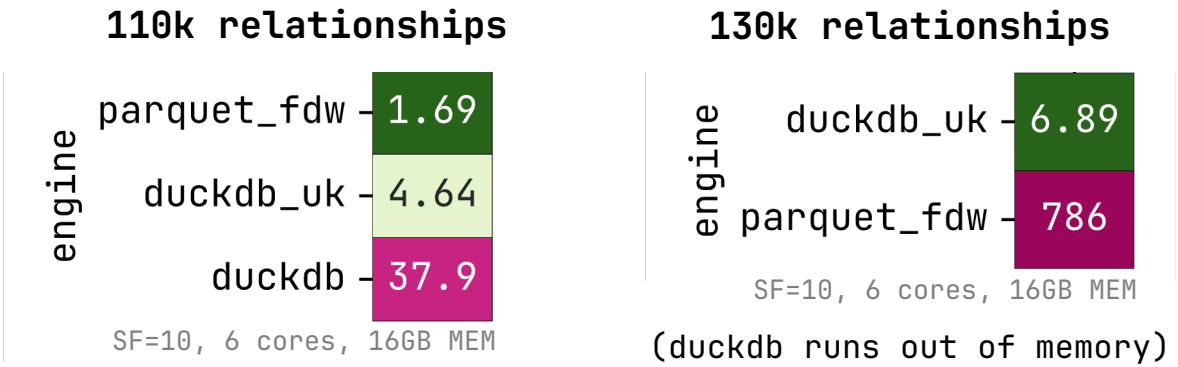

- Paths computed:
  - 110k relationships:
    - Regular CTE: 194k
    - USING KEY: 9k
  - 130k relationships:
    - Regular CTE: 74M
    - USING KEY: 19k

Summary

- Full recursive CTE support is missing in many analytical systems.
- Correct join ordering is crucial in achieving good results in graph workloads.
- DuckDB shows the overall best performance, CedarDB shows the best recursive CTE performance.
- USING KEY can help improve graph-based algorithms implemented with recursive CTEs

A SQL OLAP benchmark based on the LDBC Social Network Benchmark

Nuno Faria

INESCTEC GDC

Graph Data Council Twenty-First TUC Meeting - Boston, Sep 2026