

Lineage Capture Trade-offs: A Case Study in DuckDB

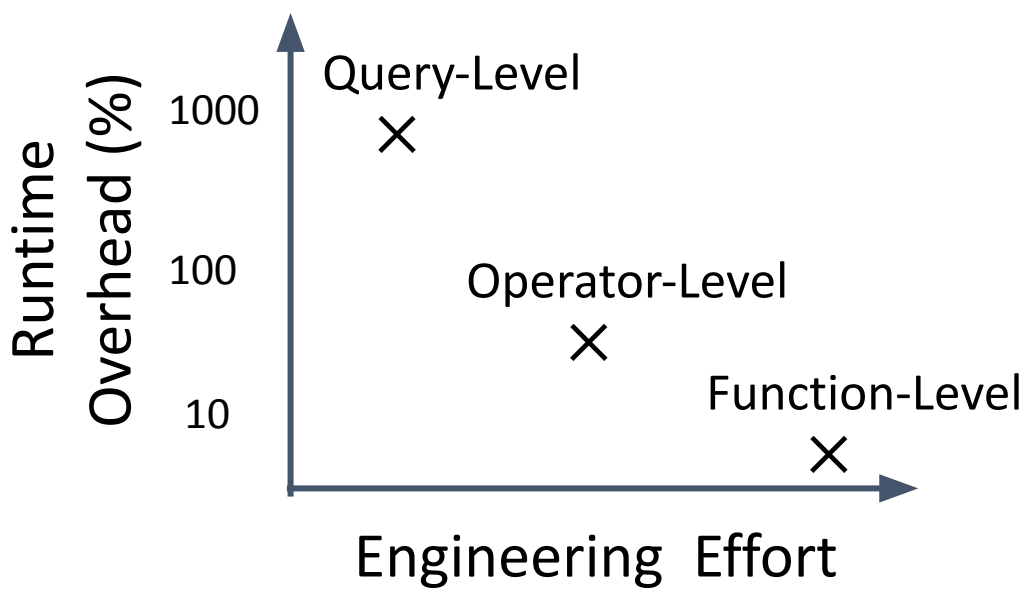
Haneen Mohammed | Columbia University, ham2156@columbia.edu
Eugene Wu | Columbia University, ewu@cs.columbia.edu


Best Lineage Capture Method in High Performance Systems?

We compare three main methods that instrument queries at different granularities

Experiment Design

- Lineage capture methods implemented in different systems, not comparable
- Implemented three methods in DuckDB for apples-to-apples comparison
- Engineering effort estimated by number of files modified





b2	cnt	A.rid	B.rid
f	2	1	1
f	2	2	1
c	2	1	2
c	2	2	2

$T1 = \gamma_{b2, count(*)}(A \times B)$
 $Q+ = T1 \bowtie \Pi_{A.rid, B.rid, b2}(A \times B)$


Query Rewrite

$Q = \gamma_{b2, count(*)}(A \times B)$

What? PERM-style query rewriting

Pros: DBMS agnostic

Cons: Logically annotates Q with prov annotations
Accumulated annotations slow down exec



b2	cnt
f	2
c	2


Plan Rewrite

$a2$ rid $b1$ $b2$ rid
 2 1 6 f 1
 3 2 6 f 1
 3 1 1 c 2
 3 2 2 c 2

What? PERM-style rewrites *per-operator* + new LM operator to strip away annotations.

Pros: Doesn't accumulate annotations during execution

Cons: Creating annotations (blue columns) still expensive at pipeline breakers
Must modify query planner



b2	cnt
f	2
c	2

FuncLib

Scatter()
Slice()
Gather()

DBMSes Rewrite

Offline

$Y+.GetData(...)$ { ...
 $Log(\{ 'Gather', idx \})$; ... }
 $Sink(...)$ { ...
 $Log(\{ 'UpdateState', idx \})$; ... }
 $Execute(...)$ { ...
 $Log(\{ 'Reference', rhs_offset, lhs_offset, count \})$; ... }


At Runtime

$Y+.GetData, [0, 1]$
 $Y+.Sink, [1, 1]$
 $x+.Exec, 0, 1, 2$
 $Y+.Sink, [0, 0]$
 $x+.Exec, 0, 0, 2$

What? Persists program variables that already encode data-movement (lineage) during execution.

Pros: Logs existing variables, avoids annotations

Cons: Must modify engine implementation



- Takeaways**
- Query-level: DBMS-agnostic but too slow
 - Operator-level: Efficient for pipelined operators and integrates cleanly with extensible query planner
 - Function-level: Faster but requires invasive DBMS changes
 - Hybrid of Function- and Operator-level may offer the best trade-off between performance and engineering effort.