

A cross-technology benchmark for incremental graph queries

Georg Hinkel, Antonio Garcia-Dominguez, René Schöne, Artur Boronat,
Massimo Tisi, Théo Le Calvar, Frederic Jouault, József Marton, Tamás Nyíri,
János Benjamin Antal, Márton Elekes, Gábor Szárnyas

Presenter: Gábor Szárnyas (CWI Amsterdam)

Software and Systems Modeling 2022 | MODELS 2022 J1 track

TTC 2018 “Social media” case

News Your solutions ▾ Programme Calls ▾ Aims and scope People History ▾



11th Transformation Tool Contest

A contest for users and developers of transformation tools.

Part of the Software Technologies: Applications and Foundations (STAF) federated conferences.

Hosted at the IRIT in Toulouse, France on Friday **29 June 2018**.

“Social media” case

Data Social network graph

Queries Score posts and comments

Changes New entities are inserted

Goal Keep query results up-to-date


Scoring posts

$$\text{Score} = 10 \times \#\text{comments} + \#\text{likes}$$

Score: 23


~~Score: 22~~

Score: 32




Scoring comments

For each comment, find connected components of users who liked the comment




Scoring comments

$$\text{Score} = \sum (\text{component size})^2$$




Scoring comments

$$\text{Score} = \Sigma (\text{component size})^2 = 1^2 + 1^2 = 2$$




Scoring comments

$$\text{Score} = \Sigma (\text{component size})^2 = 1^2 + 1^2 = 2$$




Scoring comments

$$\text{Score} = \Sigma (\text{component size})^2 = 3^2 = 9$$



Solutions



Solution	Data model	Variants
Active Operations Framework	EMF	1
ATL	EMF	2
Hawk	EMF	3
JastAdd	EMF	
Xtend	EMF	
YAMTL	EMF	3
NMF	NMF	2
Differential Dataflow	relational	1
PostgreSQL	relational	2
Neo4j	graph	2
GraphBLAS	matrix	2

Most solutions use the Eclipse Modeling Framework

21 solutions in total

DBMSs

Non-incremental query formulation

Examples of the how the initial query evaluation is formulated in:


- NMF
- Neo4j
- PostgreSQL

Note: Implementations can be quite complex – this is a “programming contest”

Scoring posts

Score = 10 × #comments + #likes

Traversing the Submission tree



NMF



```
post.Descendants()
```

Neo4j (Cypher)



```
MATCH (p:Post)
OPTIONAL MATCH (p)<-[:REPLY_OF*]-(c:Comment)
```

PostgreSQL (SQL:1999)




```
WITH RECURSIVE
comments_with_ancestors(id, ancestorid) AS (
  SELECT c.id, c.parentid AS ancestorid
  FROM comments c
  UNION
  SELECT cr.id, c.parentid AS ancestorid
  FROM comments_with_ancestors cr
  , comments c
  WHERE cr.ancestorid = c.id)
```

Scoring comments

$$\text{Score} = \sum (\text{component size})^2$$

Finding connected components of Users



NMF: Tarjan's algorithm

```
let layering = Layering<IUser>.CreateLayers(  
    comment.LikedBy,  
    u => u.Friends.Intersect(comment.LikedBy))  
let score = layering.Sum(l => Square(l.Count))
```




Neo4j: User-defined function

```
MATCH (c:Comment) WHERE (c)-[:LIKES]-(:User)  
CALL gds.wcc.stream({  
    nodeQuery:  
        "MATCH (c:Comment)-[:LIKES]-(:User)  
        WHERE id(c)='" + id(c) + "  
        RETURN id(u) AS id",  
    relationshipQuery:  
        "MATCH (u1:User)-[:FRIENDS]-(:User)  
        RETURN id(u1) AS source, id(u2) AS target",  
    validateRelationships: false  
})  
YIELD componentId  
...
```

Scoring comments

$$\text{Score} = \sum (\text{component size})^2$$


Finding connected components of Users

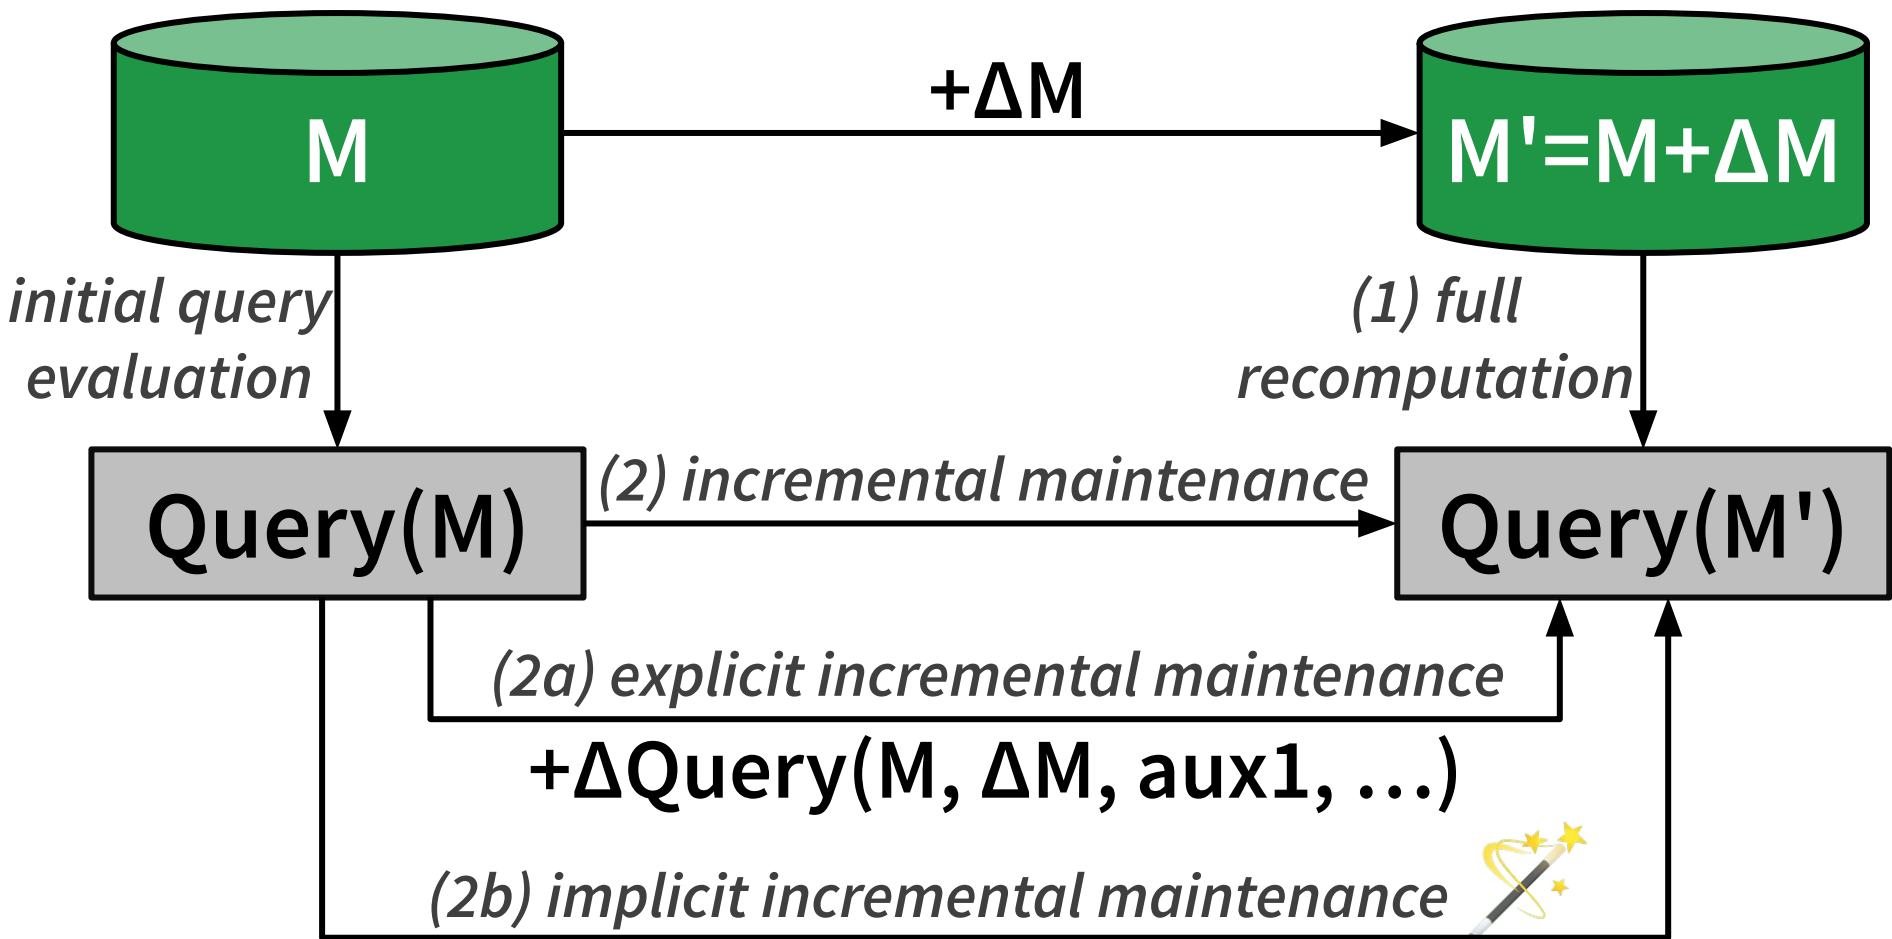


PostgreSQL: Simplified SQL query

```
WITH RECURSIVE
comment_friends(commentid, userid, user2id) AS ( ... ),
comment_friends_closed(commentid, head_userid, tail_userid) AS
  SELECT 1.commentid
        , 1.userid AS head_userid, 1.userid AS tail_userid
  FROM likes 1
  UNION
  SELECT cfc.commentid, cfc.head_userid, f.user2id AS tail_userid
  FROM comment_friends_closed cfc, comment_friends f
  WHERE cfc.tail_userid = f.userid
  AND cfc.commentid = f.commentid
), comment_components AS (
  SELECT commentid, head_userid AS userid
        , min(tail_userid) AS componentid
  FROM comment_friends_closed
  GROUP BY commentid, head_userid
), comment_component_sizes AS (
  SELECT cc.commentid, cc.componentid, count(*) AS component_size
  FROM comment_components cc
  GROUP BY cc.commentid, cc.componentid
)
SELECT c.id AS commentid
      , coalesce( sum( power(ccs.component_size, 2) ), 0) AS score
FROM comments c
  LEFT JOIN comment_component_sizes ccs ON (ccs.commentid = c.id)
GROUP BY c.id, c.ts
...
```

Incremental maintenance






Solution	Data model	Explicitly incremental	Implicitly incremental
Xtend	EMF	-	-
Hawk	EMF	+	-
PostgreSQL	relational	+	-
Neo4j	graph	+	-
GraphBLAS	matrix	+	-
Active Operations Framework	EMF	-	+
ATL	EMF	-	+
JastAdd	EMF	-	+
NMF	NMF	-	+
Differential Dataflow	relational	-	+
YAMTL	EMF	+	+

Scoring comments

$$\text{Score} = \Sigma (\text{component size})^2$$

Finding connected components of Users



Incremental evaluation

The **granularity** of the incremental maintenance has a big effect on performance:

- New “likes” edge → recalculate only for the affected single comment
- New “knows” edges → recalculate for each affected comments
- Reusing existing connected components?

Scoring comments: SQL incremental

```
1 INSERT INTO comment_friends (status,
2   commentid, userid, user2id)
3   SELECT 'B' AS status
4     , l1.commentid, f.userid,
5       f.user2id
6   FROM likes l1, likes l2
7     , friends f
8  WHERE l1.userid = f.userid
9     AND f.user2id = l2.userid
10    AND l1.commentid = l2.commentid;
```

Listing 36 Initialization phase for the Incremental PostgreSQL solution for Q2, initializing the *comment_friends* relation.

```
1 INSERT INTO comment_friends (status,
2   commentid, userid, user2id)
3   SELECT 'D' AS status
4     , l1.commentid, f.userid,
5       f.user2id
6   FROM likes_d l1, likes l2
7     , friends f
8  WHERE l1.userid = f.userid
9     AND f.user2id = l2.userid
10    AND l1.commentid = l2.commentid
11 UNION ALL
12 SELECT 'D' AS status
13     , l1.commentid, f.userid,
14       f.user2id
15   FROM likes_b l1, likes l2
16     , friends_d f
17  WHERE l1.userid = f.userid
18     AND f.user2id = l2.userid
19     AND l1.commentid = l2.commentid
20 UNION ALL
21 SELECT 'D' AS status
22     , l1.commentid, f.userid,
23       f.user2id
24   FROM likes_b l1, likes_d l2
25     , friends_b f
26  WHERE l1.userid = f.userid
27     AND f.user2id = l2.userid
28     AND l1.commentid = l2.commentid;
```

Listing 37 SQL maintenance phase for the Incremental PostgreSQL solution for Q2: updating the *comment_friends* relation.

```
1 WITH RECURSIVE
2   comment_friends_closed_init(
3     commentid, head_userid,
4     tail_userid) AS (
5     -- transitive closure
6     (reachability-only, no path is
7     recorded)
8   -- of friendship-subgraphs defined by
9   comment likes
10  -- start with the users that liked
11  a specific comment.
12  -- They are the nodes of the
13  projected users graph for a
14  comment
15  SELECT l1.commentid, l1.userid AS
16     head_userid, l1.userid AS
17     tail_userid
18  FROM likes l1
19 UNION
20 -- expand the closure with the
21 edges of the projected
22 graph,
23 -- which is stored in
24 comment_friends table
25 SELECT cfc.commentid,
26     cfc.head_userid, f.userid
27     AS tail_userid
28  FROM comment_friends_closed_init
29     cfc
30     , comment_friends f
31  WHERE cfc.tail_userid = f.userid
32     AND cfc.commentid = f.commentid
33 )
34 INSERT INTO
35   comment_friends_closed(commentid,
36   head_userid, tail_userid)
37 select commentid, head_userid,
38   tail_userid
39 from comment_friends_closed_init w
40 left join
41   comment_friends_closed q
42 using (commentid,
43   head_userid,
44   tail_userid)
45 where q.commentid IS NULL;
```

Listing 38 SQL initialization phase for the Incremental PostgreSQL solution for Q2: initializing the *comment_friends* relation's closure.

```
1 WITH comment_components AS (
2   SELECT commentid, head_userid AS
3     userid
4     , min(tail_userid) AS
5     commentid
6   FROM comment_friends_closed
7   GROUP BY commentid, head_userid
8 )
9 , comment_component_sizes AS (
10  SELECT cc.commentid,
11     cc.componentid, count(*) AS
12     component_size
13  FROM comment_components cc
14  GROUP BY cc.commentid,
15     cc.componentid
16 )
17 -- consider all comments including
18 those without likes
19 SELECT c.id AS commentid
20 , coalesce( sum(
21   power(ccs.component_size,
22     2) ), 0) AS score
23 FROM comments c left join
24   comment_component_sizes ccs
25 on (ccs.commentid = c.id)
26 GROUP BY c.id, c.ts
27 ORDER BY sum(
28   power(ccs.component_size, 2) )
29 DESC NULLS LAST
30 , c.ts DESC LIMIT 3;
```

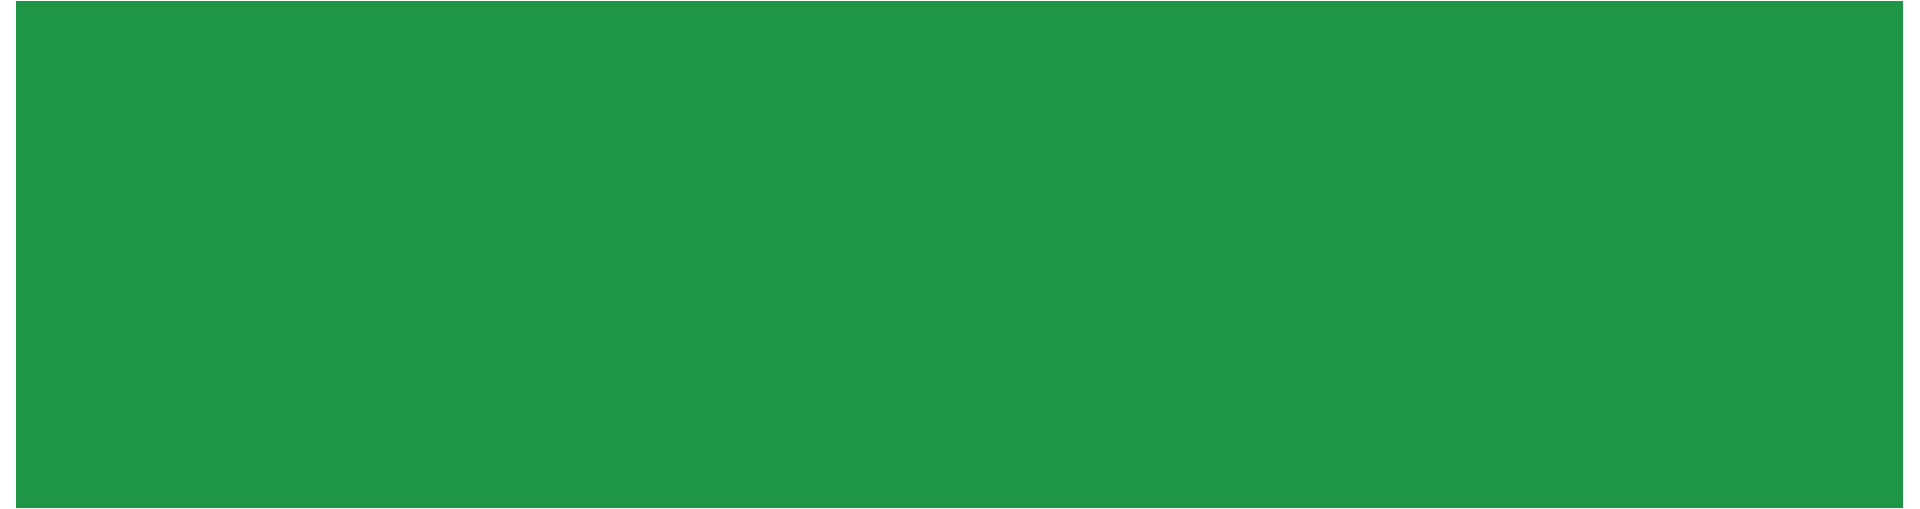
Listing 40 SQL result retrieval phase for the Incremental PostgreSQL solution for Q2.

```
1 WITH RECURSIVE -- note: though not the 1st query is
2 the recursive one, the RECURSIVE keyword
3 needs to be at the beginning
4 comment_friends_closed_stage0 AS (
5   -- in order to maintain the transitive closure in
6   comment_friends_closed
7   -- we build on the transitive closure built so
8   far and the new likes.
9   -- We need the new likes because users that liked
10  a specific comment
11  -- are the nodes of the projected users graph for
12  a comment
13  SELECT commentid, head_userid, tail_userid
14  FROM comment_friends_closed
15 UNION ALL
```


```
16 SELECT l1.commentid, l1.userid AS head_userid,
17     l1.userid AS tail_userid
18 FROM likes_d l1
19 , comment_friends_closed_stage1(commentid,
20   head_userid, tail_userid) AS (
21   -- the transitive closure computed so far
22   (reachability-only, no path is recorded)
23   -- is expanded by paths built from the new friendships
24   SELECT commentid, head_userid, tail_userid
25   FROM comment_friends_closed_stage0
26 UNION
27 SELECT cfc.commentid, cfc.head_userid, f.user2id
28   AS tail_userid
29 FROM comment_friends_closed_stage1 cfc
30   , comment_friends_d f
31 WHERE cfc.tail_userid = f.userid
32   AND cfc.commentid = f.commentid
33 )
34 , comment_friends_closed_stage2 AS (
35   -- transitive closure having the new friendships is
36   then expanded using the
37   -- previous transitive closure stage
38   SELECT distinct cfc.commentid, cfc.head_userid,
39     r.tail_userid
40   FROM comment_friends_closed_stage1 cfc
41   inner join comment_friends_closed r on
42     (cfc.tail_userid =
43     r.head_userid AND
44     cfc.commentid = r.commentid)
45 -- LEFT JOIN and WHERE ... IS NULL is
46 the antijoin
47 -- used to eliminate edges already
48 present in the previous closure
49 -- this is to prevent unnecessary
50 CONFLICTs in the INSERT
51 statement below.
52 left join comment_friends_closed s0 on
53   (cfc.commentid = s0.commentid
54   AND cfc.head_userid =
55   s0.head_userid AND
56   cfc.tail_userid =
57   s0.tail_userid)
58 WHERE s0.commentid IS NULL
59 UNION
60 SELECT commentid, head_userid, tail_userid
61 FROM comment_friends_closed_stage1
62 )
63 INSERT INTO comment_friends_closed(commentid,
64   head_userid, tail_userid)
65 select commentid, head_userid, tail_userid
66 from comment_friends_closed_stage2 w
67 left join comment_friends_closed q using
68   (commentid, head_userid,
69   tail_userid)
70 where q.commentid IS NULL
71 ON CONFLICT DO NOTHING;
```

Listing 39 SQL maintenance phase for the Incremental PostgreSQL solution for Q2: updating the *comment_friends* relation's closure.

Results and findings




Scoring posts



- AOF
- GraphBLAS Incremental
- Neo4j Incremental
- PostgreSQL Incremental
- Differential Dataflow
- JastAdd Relast Reusable Incremental
- NMF Incremental
- YAMTL Incremental

Scoring posts

Combined runtime of applying changes and re-evaluating the query



AOF

GraphBLAS Incremental

Neo4j Inc


Differential Dataflow

JastAdd Relast Reusable Incremental

NMF Incremental

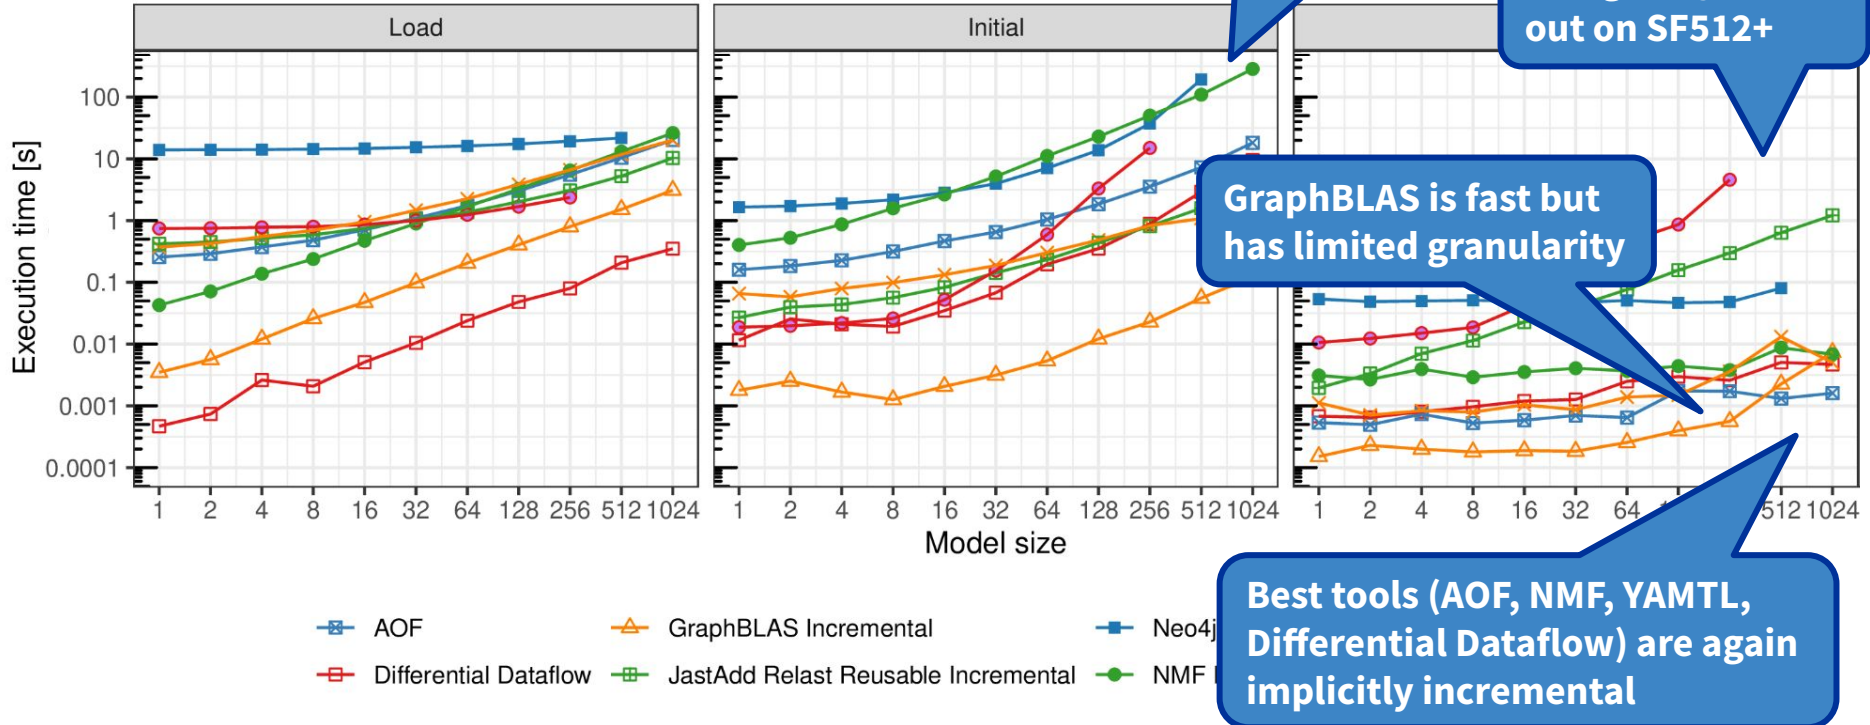
YAMTL Incremental

Scoring comments



- AOF
- GraphBLAS Incremental
- Neo4j Incremental
- PostgreSQL Incremental
- Differential Dataflow
- JastAdd Relast Reusable Incremental
- NMF Incremental
- YAMTL Incremental

Scoring comments



Findings

1. Implicitly incremental tools are superior
2. Lacklustre performance from databases
3. Parallelization is not supported by EMF tools and databases
4. User-defined functions are important
5. Fair benchmarking and reproducibility are challenging

The Linked Data Benchmark Council





The Linked Data Benchmark Council (LDBC) is a non-profit organization founded in 2012 with members from academia and industry. Its goals are:

1. Defining graph processing benchmarks
2. Facilitating fair competition
3. Accelerating the adoption of ISO GQL and SQL/PGQ

LDBC members

20 companies and organizations, including:



**ANT
GROUP**



ORACLE®
LABS

intel®



LDBC benchmarks

Social Network Benchmark:


- The data of the TTC 2018 “Social Media” case is a subset of the SNB.

The SNB has new and updated workloads:

- analytical: Business Intelligence workload v1
- transactional: Interactive workload v2

Rigorous auditing process that takes system costs (license and ownership) into account.

Conclusion



Conclusion

The TTC 2018 “Social Media” case:

- A cross-technology benchmark
- for incremental graph queries

Findings:

- Two simple graph queries can be challenging to formulate even non-incrementally
- DBMSs have performance issues for graph queries
- Explicit incremental evaluation is difficult
- Implicit incremental tools are superior
- → “retrofitted” incrementality has limited benefits

Ω

Finding: Lack of parallelization

Parallelization is paramount today: even laptop CPUs have 8–16 cores.

The initial evaluation is trivially parallelizable for both queries.

Observation:

- Only NMF, Differential Dataflow, and GraphBLAS support parallelization.
- EMF tools and databases (Neo4j, PostgreSQL) lack parallelization.

Finding: Importance of user-defined functions

Some computations are difficult to express in a declarative language, e.g. the connected components algorithm

User-defined functions (UDFs) can be used to express these computations

- Common among EMF tools – Java/Xtend code operating on the EMF model
- Database systems like Spark/Databricks and Snowflake support Java UDFs

Incremental maintenance of UDFs is difficult

Finding: Fair benchmarking is difficult

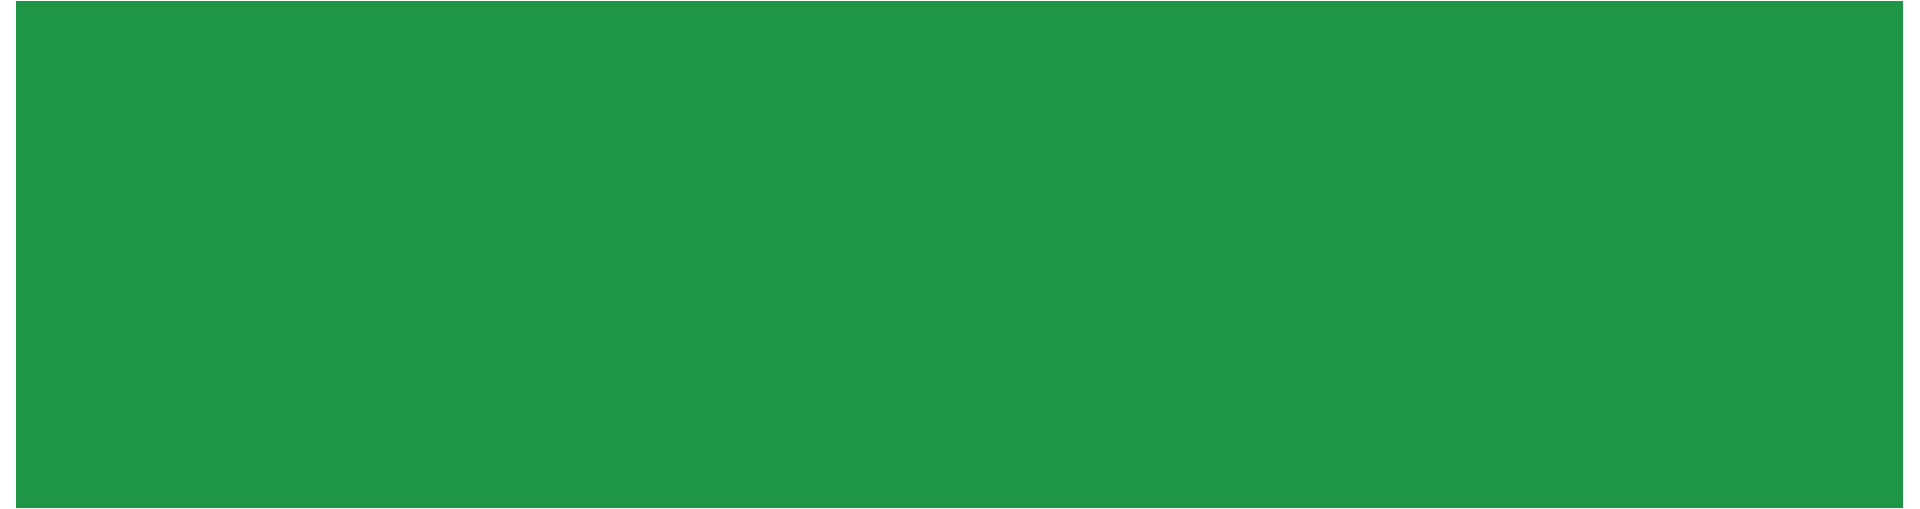
We are comparing very different systems:

- EMF tools
- Neo4j – graph DBMS
- PostgreSQL – relational DBMS
- GraphBLAS – concurrent sparse linear algebra library written in C
- Differential Dataflow – dataflow library written in Rust

Reproducibility is also difficult:

- dockerized execution
- extensive CI tests
- benchmarking in standard cloud VMs

Limitations of the benchmark



Limitations

No delete operations

- Adding them to the data generator is difficult (GRADES-NDA'20 paper)

No (unweighted) shortest path queries


- Another important graph kernel, also challenging with deletes

See examples in the next slides.

Connected components with delete operations

Scoring comments


$$\text{Score} = \sum (\text{component size})^2 = 3^2 = 9$$




Connected components with delete operations

Scoring comments


$$\text{Score} = \Sigma (\text{component size})^2 = 1^2 + 1^2 = 2$$



Shortest paths with delete operations




Shortest paths with delete operations



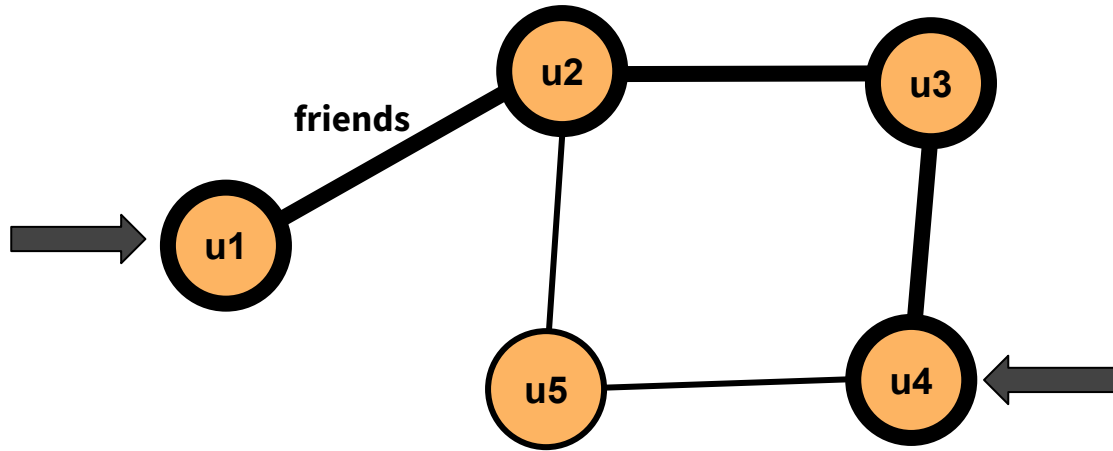
Shortest path: [u1, u5, u4]

Shortest paths with delete operations




Shortest path: ?

Shortest paths with delete operations



Shortest path: [u1, u2, u3, u4]

Shortest paths with delete operations



Shortest path: [u1, u2, u3, u4]
[u1, u2, u5, u4]

Ideas for incremental evaluation //
Future work



Ideas for incremental evaluation

The ideas in the following slides could work **if all inserts are added one-by-one** and there are not too many inserts. (There aren't, see the table with the [model sizes](#).)

IIRC none of the solutions in the paper used this: they all went for a bulk insertion followed by a single recomputation step.

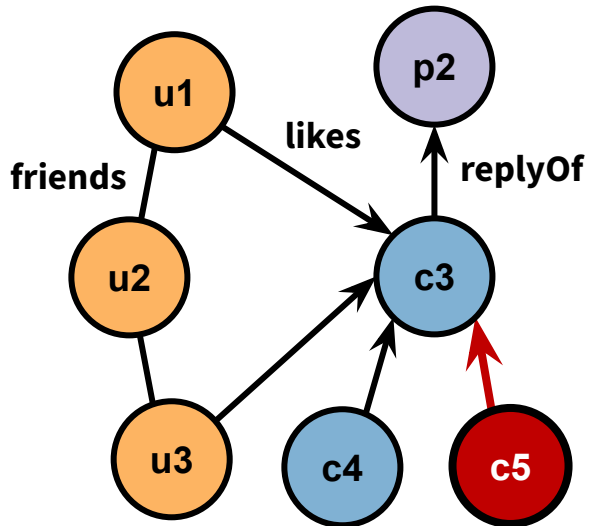
With a client-server setup, doing operations in bulk likely makes sense. Performing an individual maintenance operation per insert is likely expensive. With a read-oriented system (e.g. column store), it makes sense to perform the inserts in bulk.

Still, it would be interesting to give this a go with an embeddable database (e.g. DuckDB, Neo4j) or a system which provides an option to write stored procedure (like Oracle's PL/SQL).

Scoring posts

Score = 10 × #comments + #likes

Traversing the Submission tree



Trick: For each Comment, store its root Post. When inserting a new child Comment, it should get its parent's root Post.

This works because the subgraph is a tree and there are no cut-and-link operations.


Trick: Upon adding a new “likes” or a new “friends” edge, connected components can only be merged together.

This works because there are no delete operations.

Scoring comments

$$\text{Score} = \Sigma (\text{component size})^2$$

Finding connected components of Users




Trick: Upon adding a new “likes” or a new “friends” edge, connected components can only be merged together.

This works because there are no delete operations.

Scoring comments

$$\text{Score} = \Sigma (\text{component size})^2$$

Finding connected components of Users

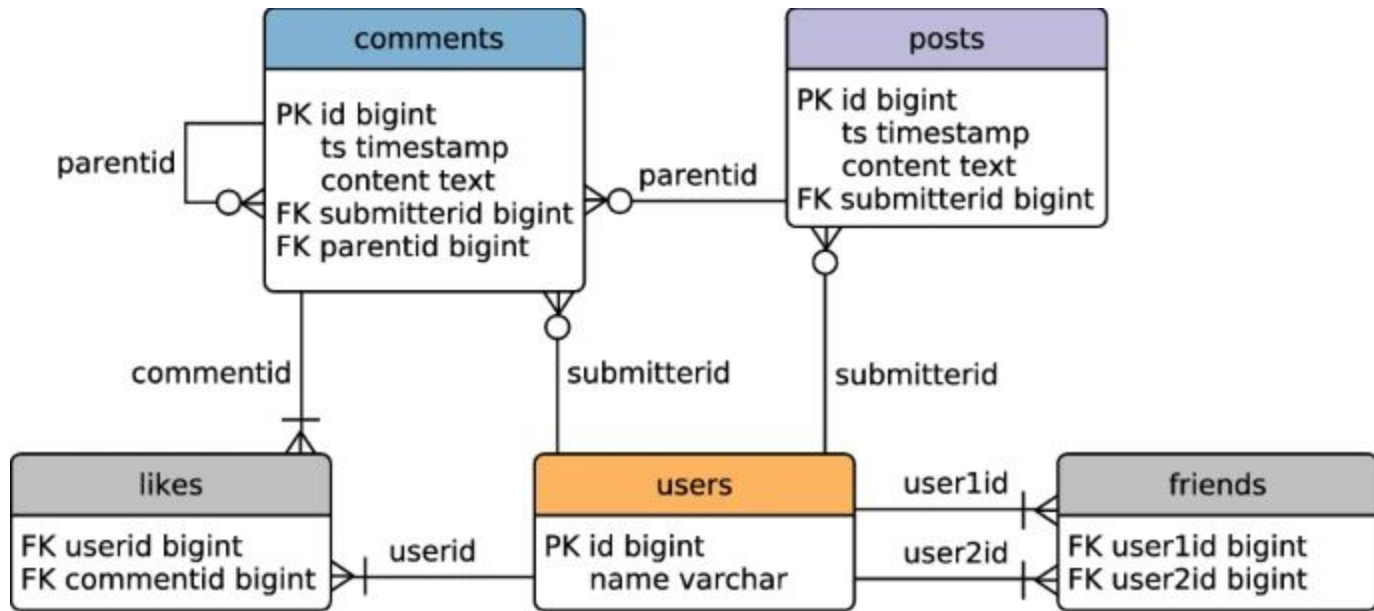


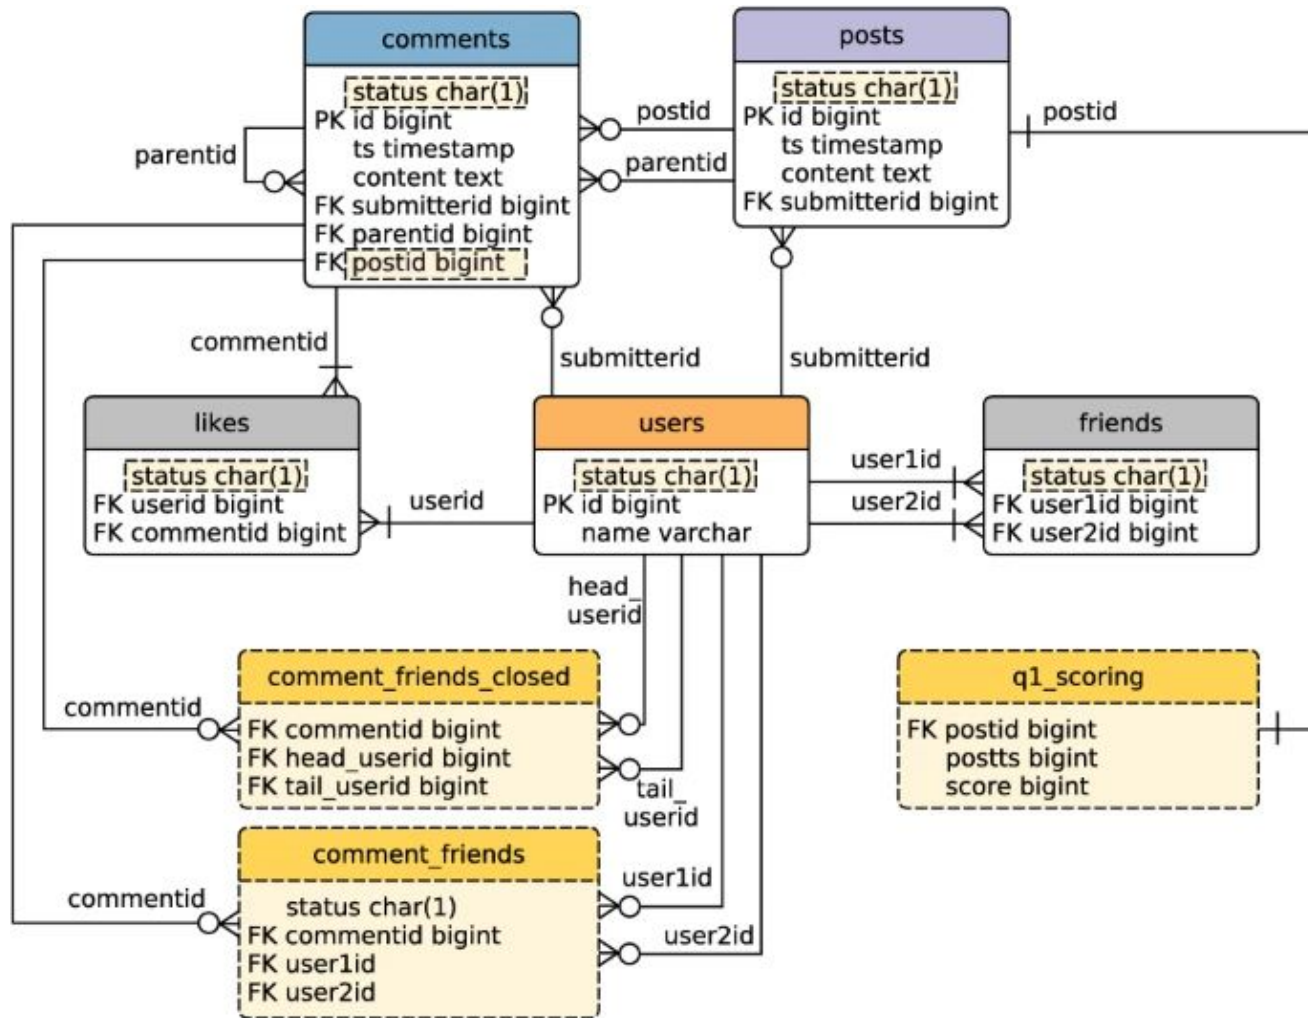
Potential extensions to the slide deck

- ~~concrete result slides~~
- ~~details on SQL/PGQ~~ ~~WONTFIX~~
- ~~model sizes~~
- ~~incremental query formulation~~
- interesting findings
- more info on concrete tools
 - mention of DD & videos
 - mention of GraphBLAS
 - Hawk, NMF, YAMTL, etc.
- anything on DuckDB/DuckPGQ as a potential tool for Q1
- ~~complaining that most MDE tools are single-threaded~~
- incremental tricks explained...

Incremental query formulation







Incremental view maintenance

Categories:


- **Non-incremental:** query is recomputed each time
- **Implicitly implemental:** the maintenance is done automatically by the system
- **Explicitly incremental:** the query developer manually incrementalizes the query
(poor man's view maintenance, can be retrofitted to existing systems)

Studied in depth in database research

...but most research focused on equijoins

...maybe anti- and outer joins

Transitive reachability (tree queries, connected components, etc.) are less studied.



■ AOF
 ▲ GraphBLAS Incremental
 ■ Neo4j Incremental
 ● PostgreSQL Incremental
■ Differential Dataflow
 ■ JastAdd Relast Reusable Incremental
 ● NMF Incremental
 × YAMTL Incremental

All tools

- (1) AOF
- (2) ATL (3) + Incremental
- (4) Differential Dataflow
- (5) GraphBLAS (6) + Incremental
- (7) Hawk (8) + IU (9) + IUQ
- (10) JastAdd
- (11) Neo4j (12) + Incremental
- (13) NMF (14) + Incremental
- (15) PostgreSQL (16) + Incremental
- (17) Xtend
- (18) YAMTL (19) + II (20) + EI

Most tools use the EMF data model

(1) Active Operations Framework (AOF)

(2) ATL (3) + Incremental

(7) Hawk (8) + IU (9) + IUQ

(10) JastAdd

(17) Xtend

(18) YAMTL (19) + II (20) + EI

NMF:

(13) NMF (14) + Incremental

Relational:

(4) Differential Dataflow

(15) PostgreSQL (16) + Incremental

Property graph:

(11) Neo4j (12) + Incremental

Matrix:

(5) GraphBLAS (6) + Incremental

Draft

a very small benchmark suite, just two queries and a few transformations

already highlights numerous usability and performance characteristics of systems !!

e.g. why don't MDE tools use relational DBMSs

discuss the two queries briefly

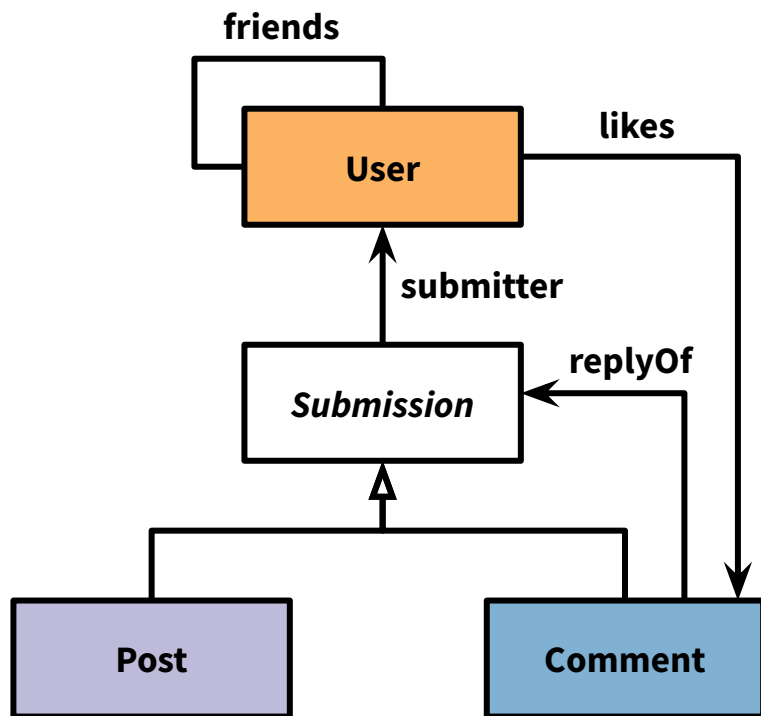
present a few solutions (e.g. Postgres/SQL; Neo4j/Cypher; MDE tools; differential dataflow, refer to Frank McSherry's video)

FMS videos


Live coding differential dataflow:

- https://www.youtube.com/watch?v=W6TKxS_pWr0
- <https://www.youtube.com/watch?v=83rG471bmw8>
- <https://www.youtube.com/watch?v=uZ23MnpujNA>

Schema



Instance



Differential dataflow: CC computation

```
likes // node label comment
.filter(|_| false)
.map(|(user, comm)| ((user.clone(), comm), user))
.iterate(|labels| {

    let knows = knows.enter(&labels.scope());
    let likes = likes.enter(&labels.scope());

    labels
        .map(|((node, comment), label)| (node, (label, comment)))
        .join_map(&knows, |_node, (label, comment), dest| ((dest.clone(), comment.clone()), label.clone()))
        .concat(&likes.map(|(user, comm)| ((user.clone(), comm), user)))
        .reduce(|_key, input, output| {
            // only produce output, if `input` contains `_key.0`
            if input.iter().any(|(label, _wgt)| *label == &_key.0) {
                output.push((input[0].0.clone(), 1));
            }
        })
});
```

Table 2 Classification of approaches for Q2, sorted by tool and incrementality

Solution	Incremental conn. components	Algorithm	Sorting
AOF	⊗	Breadth-first traversal	Incremental
ATL	○	Depth-first traversal	Full
ATL incremental	⊗	Breadth-first traversal	Incremental
Differential Dataflow	⊗	Fixed-point label propagation	Full
GraphBLAS	○	FastSV	Offline top- x
GraphBLAS incremental	⊗	FastSV on overestimation + merge	Online top- x
Hawk	○	Tarjan	Full
Hawk (IU)	○	Tarjan	Full
Hawk (IUQ)	○	Tarjan	Online top- x
JastAdd	⊗	Depth-first/Kosaraju	Offline top- x
Neo4j	○	Union-find variant	Offline top- x
Neo4j incremental	⊗	Breadth-first traversal	Incremental
NMF reference	○	Tarjan	Full
NMF incremental	⊗	Edge changes	Incremental
PostgreSQL	○	Breadth-first traversal	Online top- x
PostgreSQL incremental	⊗	Overestimation + breadth-first traversal	Online top- x
Xtend	○	Tarjan	Online top- x
YAMTL-B	○	Weighted quick-union-find with path compression	Online top- x
YAMTL-II	⊗	Weighted quick-union-find with path compression	Online top- x
YAMTL-EI	⊗	Weighted quick-union-find with path compression	Online top- x

Notation—⊗ yes; ⊗ to some extent; ○ no. Overestimation means that all connected components are re-computed that might be affected by the changes

Table 1 Classification of approaches for Q1, ordered by solution name and support for incrementality

Solution	Incrementality	Sorting
AOF	⊗	Incremental
ATL	○	Full
ATL incremental	⊗	Incremental
Differential Dataflow	⊗	Full
GraphBLAS	○	Offline top- x
GraphBLAS incremental	⊗	Offline top- x
Hawk	∅	Full
Hawk (IU)	∅	Full
Hawk (IUQ)	∅	Online top- x
JastAdd	∅	Offline top- x
Neo4j	○	Offline top- x
Neo4j incremental	⊗	Incremental
NMF reference	○	Full
NMF incremental	⊗	Incremental
PostgreSQL	○	Online top- x
PostgreSQL incremental	⊗	Online top- x
Xtend	○	Online top- x
YAMTL-B	○	Online top- x
YAMTL-II	∅	Online top- x
YAMTL-EI	⊗	Online top- x

Notation: ⊗ yes; ∅ to some extent; ○ no

Table 4 Comparison of the tools used in the paper

Tool	Version	Data model	Engine	Solution	Decl.	Batch	Implicit	Explicit	DB	MV	Parallel
AOF	v201806	EMF	Java 8	Xtend	⊘	○	⊗	○	○	○	○
ATL	3.8.0	EMF	Java 8	ATL	⊗	⊗	○	○	○	○	○
ATL Incremental	v201904	EMF	Xtend	ATL/AOF	⊗	○	⊗	○	○	○	○
Differential Dataflow	0.11.0	relations	Rust	Rust	○	○	⊗	○	○	○	⊗
GraphBLAS	4.0.3	matrices	C	C++	○	⊗	○	⊗	○	○	⊗
Hawk	2.1.0	EMF	Java 8	EOL	⊘	⊗	⊘	⊗	⊗	⊘	○
JastAdd	2.3.5	EMF	Java 11	Java 11	○	⊗	⊗	○	○	○	○
Neo4j	4.2.4	property graph	Java 11	Java 11	⊗	⊗	○	⊗	⊗	⊗	○
NMF	2.0.169	NMF	C#	C#	⊗	⊗	⊗	○	○	○	⊗
PostgreSQL	12.4	relations	C	Java 11/SQL	⊗	⊗	○	⊗	⊗	⊗	○
Xtend	2.20.0	EMF	Java 8	Xtend	⊗	⊗	○	○	○	○	⊗
YAMTL	0.1.5	EMF	Java 11	Xtend	⊘	⊗	⊗	⊗	○	○	○

Data Model: the data model exposed to the user. *Engine*: programming language used to implement the engine (model transformation engine, database query engine, etc.), *Solution*: programming language and query language (if applicable) used to implement the solution, *Decl.*: the solution specified the queries using a declarative query language, *Batch*: only batch mode is supported, *Implicit*: implicit incrementalization is supported, *Explicit*: a solution with explicit incrementalization was implemented, *DB*: database-backed, i.e. the tool persists the model on disk after each transaction, *MV*: materialized views, *Parallel*: parallelization is supported. *Notation*—⊗ yes; ⊘ to some extent; ○ no; ⊕ yes, using Java 8 streams

Model sizes for each scale factor

Table 7 Model sizes for each scale factor: number of nodes and edges, number of changes

Type\scale factor	1	2	4	8	16	32	64	128	256	512	1024
Comments	640	1064	2315	5056	9220	18,872	39,212	76,735	148,470	273,418	540,905
Posts	554	889	1845	2270	5518	10,929	18,083	37,228	74,668	167,299	314,510
Users	80	118	190	204	394	595	781	1158	1678	2606	3699
Total number of nodes	1274	2071	4350	7530	15,132	30,396	58,076	115,121	224,816	443,323	859,114
friends	53	102	262	298	904	1827	2752	5695	11,118	24,387	45,386
replyTo	640	1064	2315	5056	9220	18,872	39,212	76,735	148,470	273,418	540,905
likes	6	24	66	129	572	1598	4770	13,374	36,815	102,276	268,432
submitter	1194	1953	4160	7326	14,738	29,801	57,295	113,963	223,138	440,717	855,415
Total number of edges	2533	4207	9118	17,865	34,654	70,970	143,241	286,502	568,011	1,114,216	2,251,043
Total number of changes	67	120	132	104	110	117	68	86	45	112	74