

# The case against specialized graph engines

*Jing Fan, Adalbert Gerald Soosai Raj, and Jignesh M. Patel*  
*University of Wisconsin – Madison*

# Motivation

- Graph analytics is now common
- Response = new specialized graph engines

## “One Size Fits All”: An Idea Whose Time Has Come and Gone


Michael Stonebraker  
*Computer Science and Artificial  
Intelligence Laboratory, M.I.T., and  
StreamBase Systems, Inc.*  
stonebraker@csail.mit.edu

Uğur Çetintemel  
*Department of Computer Science  
Brown University, and  
StreamBase Systems, Inc.*  
ugur@cs.brown.edu



# Motivation


- Graph analytics is now common
- Response = new specialized graph engines



# Motivation

- Graph analytics is now common
- Response = new specialized graph engines

Google  
Pregel



Stanford GPS

GraphLab  
Unleash Data Science™

Question: Is graph processing that different from other types of data processing?

**Our Answer: No. Can be subsumed by “traditional” relational processing**


# What is appealing about these new engines?

Vertex-  
centric API

Easy to write  
graph  
programs

Higher  
programmer  
productivity


# Graph API: Giraph



## Vertex Centric:


```
do {  
  foreach vertex in the graph {  
    receive_messages();  
    mutate_vertex_value();  
    if (send_to_neighbors()) {  
      send_messages_to_neighbors();  
    }  
  }  
} until (has_converged() || reached_limit())
```

# Example: Shortest path




## Computation & Communication Pattern

**Iteration 1**




# Example: Shortest path




## Computation & Communication Pattern

**Iteration 2**






# Example: Shortest path




## Computation & Communication Pattern



**Iteration 3**

# GraphLab



1. **Gather** values (from neighbors)
2. **Apply** updates to local state
3. **Scatter** signals to your neighbors

# What is appealing about these new engines?

Vertex-  
centric API

Easy to write  
graph  
programs

Higher  
programmer  
productivity

## But ...

- Can we build a similar vertex-centric simple API?
- ... and then map it to SQL, with good performance




## The GRAIL API

### Advantages:

- Already have SQL in the enterprise stack
- Huge advantage to “one size fits many”
  - $O(N^2)$  headaches when maintaining  $N$  specialized systems
  - Economies of scale

# Example: Shortest path



**Input Graph**

**Vertex**

id	val
A	$\infty$
B	$\infty$
C	$\infty$
D	$\infty$

**Edge**

src	dest	val
A	B	1
A	C	2
B	D	2
C	D	3

**next**

id	val
A	0
B	$\infty$
C	$\infty$
D	$\infty$

**message**

id	val
B	1
C	2

**Iteration 1**

**next**

id	val
A	0
B	1
C	2
D	$\infty$

**message**

id	val
D	3
D	5

**Iteration 2**

**next**

id	val
A	0
B	1
C	2
D	3

**message**

id	val

**Iteration 3**

# The Grail API

T-SQL Code

```
1 DECLARE @flag int;
2 SET @flag = 1;
3 SELECT vertex.id, 2147483647 AS val
4 INTO next
5 FROM vertex;
6 CREATE TABLE message(
7   id int,
8   val int
9 );
10 INSERT INTO message values(1,0);
11 WHILE (@flag != 0)
12 BEGIN
13   SELECT message.id AS id, MIN(message.val) AS val
14   FROM message
15   GROUP BY message.id;
16 DROP TABLE message;
17
18 SELECT cur.id AS id, cur.val AS val
19 INTO update
20 FROM cur, next
21 WHERE cur.id = next.id AND cur.val < next.val;
22 SET next.val = update.val
23 WHERE next.id = update.id;
24 SELECT edge.dest AS id, update.val + 1 AS val
25 INTO message
26 FROM update, edge
27 WHERE edge.src = update.id;
28 DROP TABLE cur;
29 DROP TABLE update;
30
31
32
33
```

**Initialize**

**Initialize the message table**

**Aggregate the messages**

**Update and generate messages for the next iteration**

**Stop when no new msgs.**

```
1 VertexValType: INT
2 MessageValType: INT
3 InitiateVal : INT_MAX
4 InitialMessage : (1, 0)
5 CombineMessage: MIN(message)
6 UpdateAndSend: update=cur.val<getVal()
7     if (update) {
8         setVal(cur.val)
9         send(out, cur.val+1)
10    }
11 End: NO_MESSAGE
```

# T-SQL Code

```
1 DECLARE @flag int;
2 SET @flag = 1;
3 SELECT vertex.id, 2147483647 AS val
4 INTO next
5 FROM vertex;
6 CREATE TABLE message(
7     id int,
8     val int
9 );
10 INSERT INTO message values(1,0);
11 WHILE (@flag != 0)
12 BEGIN
13     SELECT message.id AS id, MIN(message.val) AS val
14     INTO cur
15     FROM message
16     GROUP BY message.id;
17 DROP TABLE message;
18 SELECT cur.id AS id, cur.val AS val
19 INTO update
20 FROM cur, next
21 WHERE cur.id = next.id AND cur.val < next.val;
22 UPDATE next
23 SET next.val = update.val
24 FROM update, next
25 WHERE next.id = update.id;
26 SELECT edge.dest AS id, update.val + 1 AS val
27 INTO message
28 FROM update, edge
29 WHERE edge.src = update.id;
30 DROP TABLE cur;
31 DROP TABLE update;
32 SELECT @flag = COUNT(*) FROM message;
33 END
```

Initialize

Initialize the message table

Aggregate the messages

Create an update table and only consider updated vertices

Update the next table

Generate the message table for the next iteration

Stop when there are no new messages

Vertex Centric	Relational Algebra
Receive messages	$cur \leftarrow \gamma_{id, F_0(val)}(message)$
Mutate value	$next \xleftarrow{u} \pi_{next.id, F_1(other.val)} other \bowtie_{id} next$
Send messages	$\pi_{edge.B, F_2(other.val, edge.val)} other \bowtie_{other.id=edge.A} edge$

Aggregate function  
(can be a UDAF)

Scalar computation  
(can be a UDF)

Scalar computation  
(can be a UDF)

Join attributes  
control the direction

For single source  
shortest path

min


sum

identity

Outgoing edges



# Grail: Implementation and Evaluation



## Test Machine (single node)

- Dual 1.8GHz Xeon E2450L
- 96GB of main memory

## Compare with

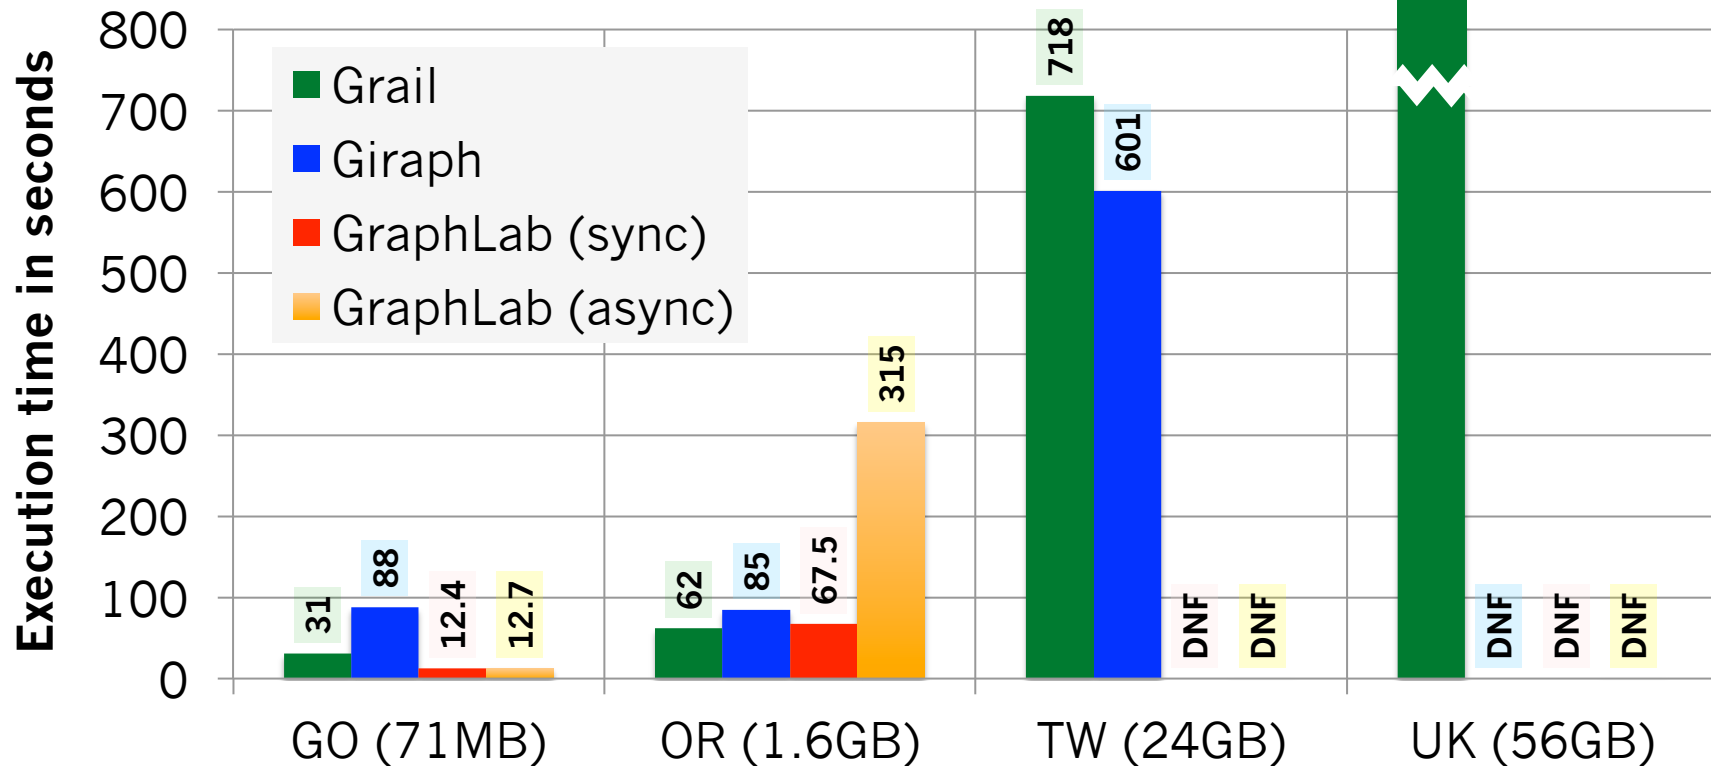
- Giraph (v.1.1.0)
- GraphLab (v 2.2): sync and async

Dataset	#nodes	#edges	size
web-google (GO)	9K	5M	71MB
com-Orkut (OR)	3M	117M	1.6GB
Twitter-10 (TW)	41.6M	1.5B	24GB
uk-2007-05 (UK)	100M	3.3B	56GB

## Queries

- Single source shortest-path
- Page Rank
- Weakly connect components

# Results: Single Source Shortest Path



Grail is slower than GraphLab for the smallest datasets,  
... but catches up as the dataset size grows,  
... and can handle the largest datasets, while the other systems fail

# Summary: Graph Analytics on RDBMS



Simple API (Grail) addresses the programmer productivity issue



Produces far more robust and deployable solutions than specialized graph engines



Interesting physical schema design and optimization issues

# The general case against GraphDB Inc.




Figure 1: *Fraction of jobs per application type.*




Figure 2: *Fraction of users per application type.*

# Thanks!

*Microsoft*  
**GRAY SYSTEMS LAB**



David DeWitt



Jae Young Do



Alan Halverson



Ian Rae