NEDB Day 2012

A Parallel System for Efficiently Managing Large Graphs

Jeong-Hyon Hwang, Sean Spillane, Daniel Bokser,

Daniel Kemp, Jayadevan Vijayan, Jeremy Birnbaum



State University of New York - Albany

• Transportation

5:00 AM - 6:00 AM



• Transportation

5:00 AM - 6:00 AM



• Transportation

5:00 AM - 6:00 AM





• Transportation

5:00 AM - 6:00 AM





6:00 AM - 7:00 AM

• Transportation

5:00 AM - 6:00 AM



6:00 AM - 7:00 AM





• Transportation

5:00 AM - 6:00 AM





7:00 AM - 8:00 AM



• Transportation







• Transportation



NEDB Day 2012

Saturday, July 6, 2013

• Transportation



- Social and Political Studies / Marketing / National Security
- Epidemic Simulations and Analysis

	inconvenient	inefficient
RDBMSs	no SQL support for graphs	no built-in graph traversal capabilities
Graph Systems	coding for each type of computation	only one graph at a time

	inconvenient	inefficient
RDBMSs	no SQL support for graphs	no built-in graph traversal capabilities
Graph Systems	coding for each type of computation	only one graph at a time

Example: shortest distance in Pregel



	inconvenient	inefficient
RDBMSs	no SQL support for graphs	no built-in graph traversal capabilities
Graph Systems	coding for each type of computation	only one graph at a time

Goal

- convenient and efficient queries on graphs using a server cluster

• Key Features

- succinct graph query language
- distributed, deduplicate storage of graphs
- parallel query processing that shares computations across graphs

Goal

- convenient and efficient queries on graphs using a server cluster

Key Features

- succinct graph query language
- distributed, deduplicate storage of graphs
- parallel query processing that shares computations across graphs



Goal

- convenient and efficient queries on graphs using a server cluster

Key Features

- succinct graph query language
- distributed, deduplicate storage of graphs
- parallel query processing that shares computations across graphs



Goal

- convenient and efficient queries on graphs using a server cluster

Key Features

- succinct graph query language
- distributed, deduplicate storage of graphs
- parallel query processing that shares computations across graphs



Goal

- convenient and efficient queries on graphs using a server cluster

Key Features

- succinct graph query language
- distributed, deduplicate storage of graphs
- parallel query processing that shares computations across graphs



Goal

- convenient and efficient queries on graphs using a server cluster

Key Features

- succinct graph query language
- distributed, deduplicate storage of graphs
- parallel query processing that shares computations across graphs



Goal

- convenient and efficient queries on graphs using a server cluster

Key Features

- succinct graph query language
- distributed, deduplicate storage of graphs
- parallel query processing that shares computations across graphs



Goal

- convenient and efficient queries on graphs using a server cluster

Key Features

- succinct graph query language
- distributed, deduplicate storage of graphs
- parallel query processing that shares computations across graphs

































































Example: average degree of each graph in {G₁,G₂}

















query time (sec)

• Data

- 10 complete binary trees
- i'th tree = (i-1)'th tree

+

new 25K vertices (and edges)

- Query
 - single-source shortest paths

Phoebus (all graphs) Phoebus (last graph) 6000 G* (all graphs) G* (last graph) 4500 3000 1500 0 2 5 8 3 4 9 7 10

number of graphs

• Data

- 10 complete binary trees
- i'th tree = (i-1)'th tree

+

new 25K vertices (and edges)

- Query
 - single-source shortest paths



• Data

- 10 complete binary trees
- i'th tree = (i-1)'th tree

╋

new 25K vertices (and edges)

- Query
 - single-source shortest paths



• Data

- 10 complete binary trees
- i'th tree = (i-1)'th tree

╋

new 25K vertices (and edges)

- Query
 - single-source shortest paths



• Data

- 10 complete binary trees
- i'th tree = (i-1)'th tree

╋

new 25K vertices (and edges)

- Query
 - single-source shortest paths



Scale-up Test

- Platform
 - 64 core server cluster
- Data
 - I complete binary tree
 - I billion vertices (edges)
 - 64 million vertices / server
- Queries
 - vertex degree
 - single-source shortest paths



Scale-up Test

- Platform
 - 64 core server cluster
- Data
 - I complete binary tree
 - I billion vertices (edges)
 - 64 million vertices / server
- Queries
 - vertex degree
 - single-source shortest paths



number of cores (relative data size)

Scale-up Test

- Platform
 - 64 core server cluster
- Data
 - I complete binary tree
 - I billion vertices (edges)
 - 64 million vertices / server
- Queries
 - vertex degree
 - single-source shortest paths



number of cores (relative data size)

Conclusion

- Convenient and efficient queries on graphs
 - distributed, deduplicate storage of graphs
 - parallel query processing that share computations across graphs
- Supported by NSF CAREER award IIS-1149372
- 30K lines of Java code
- Ist code release in August 2012
- Future work
 - disk optimization / graph distribution / query optimization / high availability

NEDB Day 2012

Thank you!