The Graphalytics Ecosystem

From Competitions to Performance Analysis

June 20, 2018


Ahmed Musaafir


Tim Hegeman


Wing Lung Ngai


Dr. Alexandru Uta


Prof.dr.ir. Alexandru Iosup

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft Delft University of Technology

atlarge-research.com
graphalytics.org

The Data Deluge: Large-scale Graphs

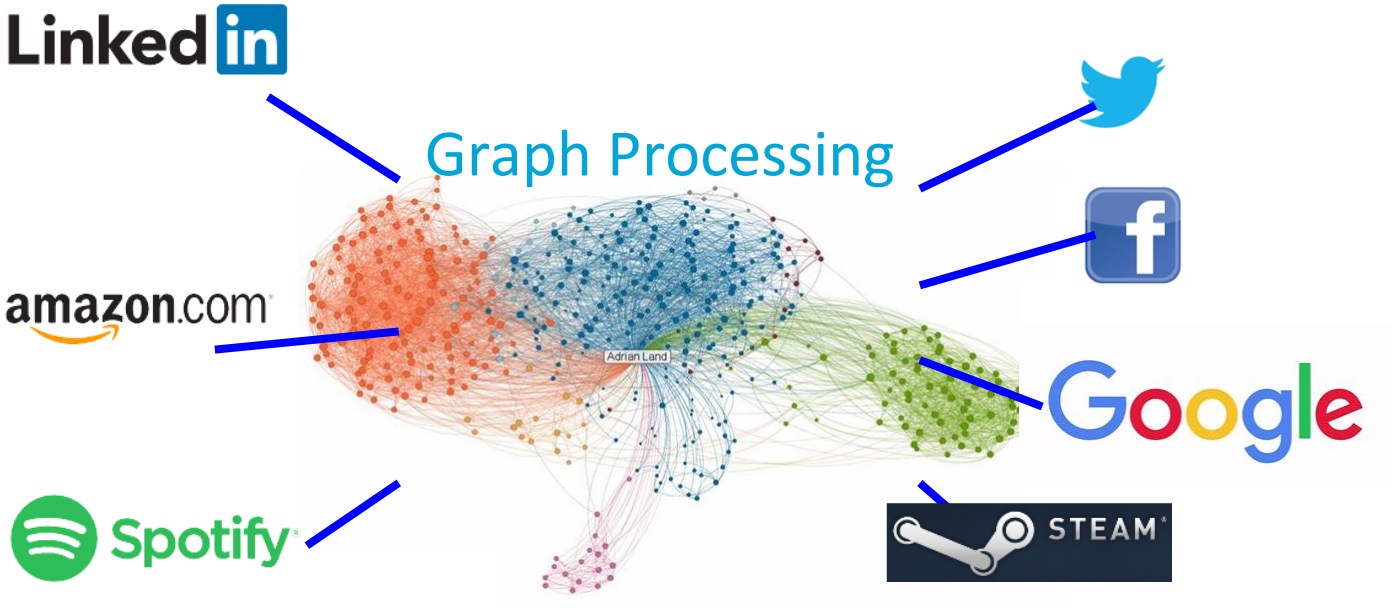

The Data Deluge: Large-scale Graphs

Predicting or recommending new relationships (friends-of-friends, product recommendations).

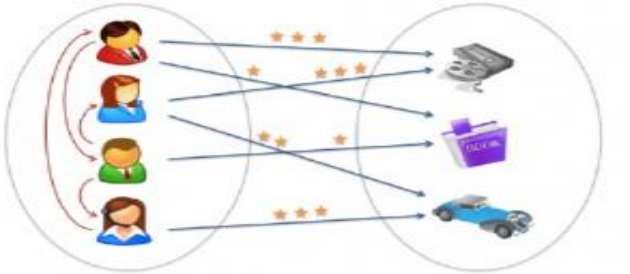

Navigation systems

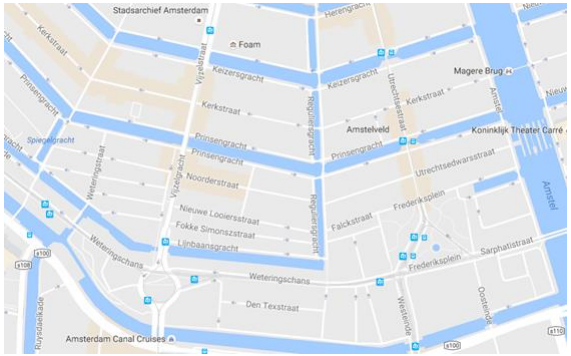

The Data Deluge: Large-scale Graphs

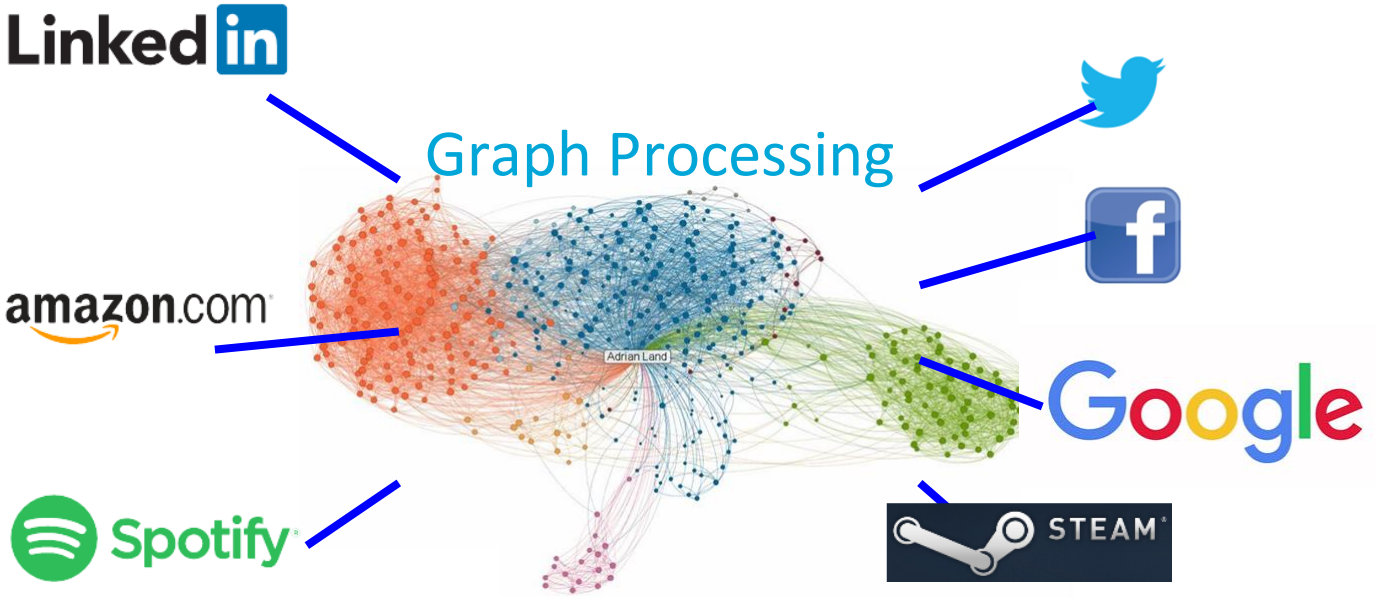

The Data Deluge: Large-scale Graphs

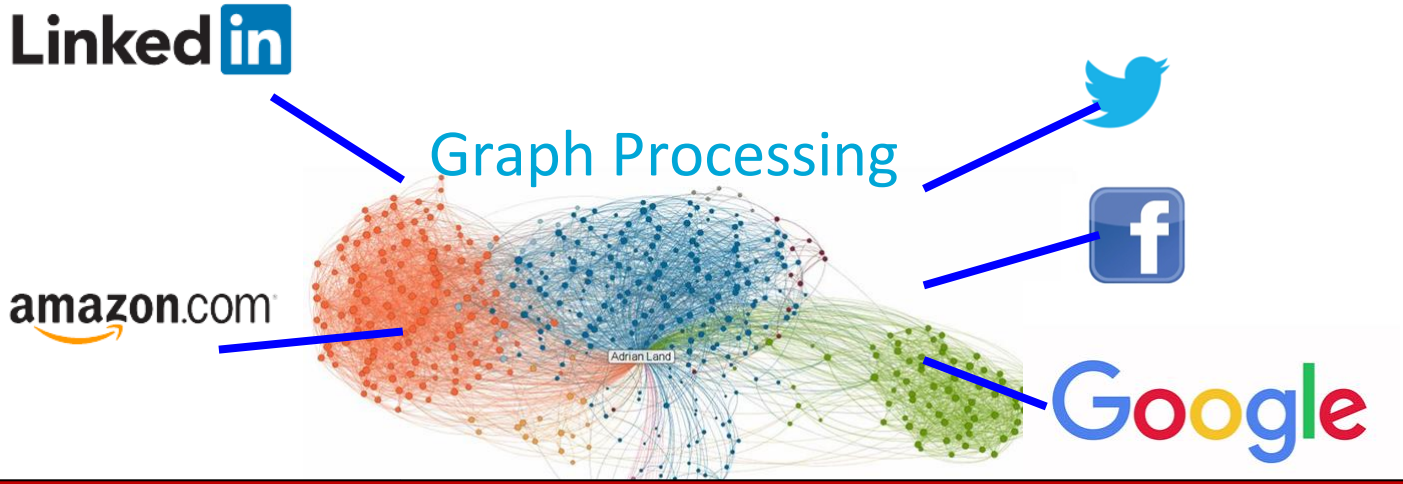

Analyse large graphs

Graph Processing Platforms

Intel Graphmat

Gemini

IBM System G

JoyGraph

HyGraph

Graph Processing Platforms

Q1: Which platforms perform well?

IBM System G

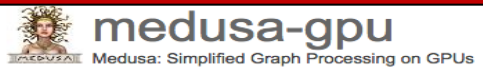

JoyGraph

HyGraph

Graph Processing Platforms

Understanding Graph Processing Performance

Two dimensions for understanding performance:

Breadth: comparison across diverse platforms, algorithms, datasets.

Answers Q1: which platforms performs well?

Depth: performance analysis of individual jobs.

Answers Q2: why?

The Graphalytics Ecosystem

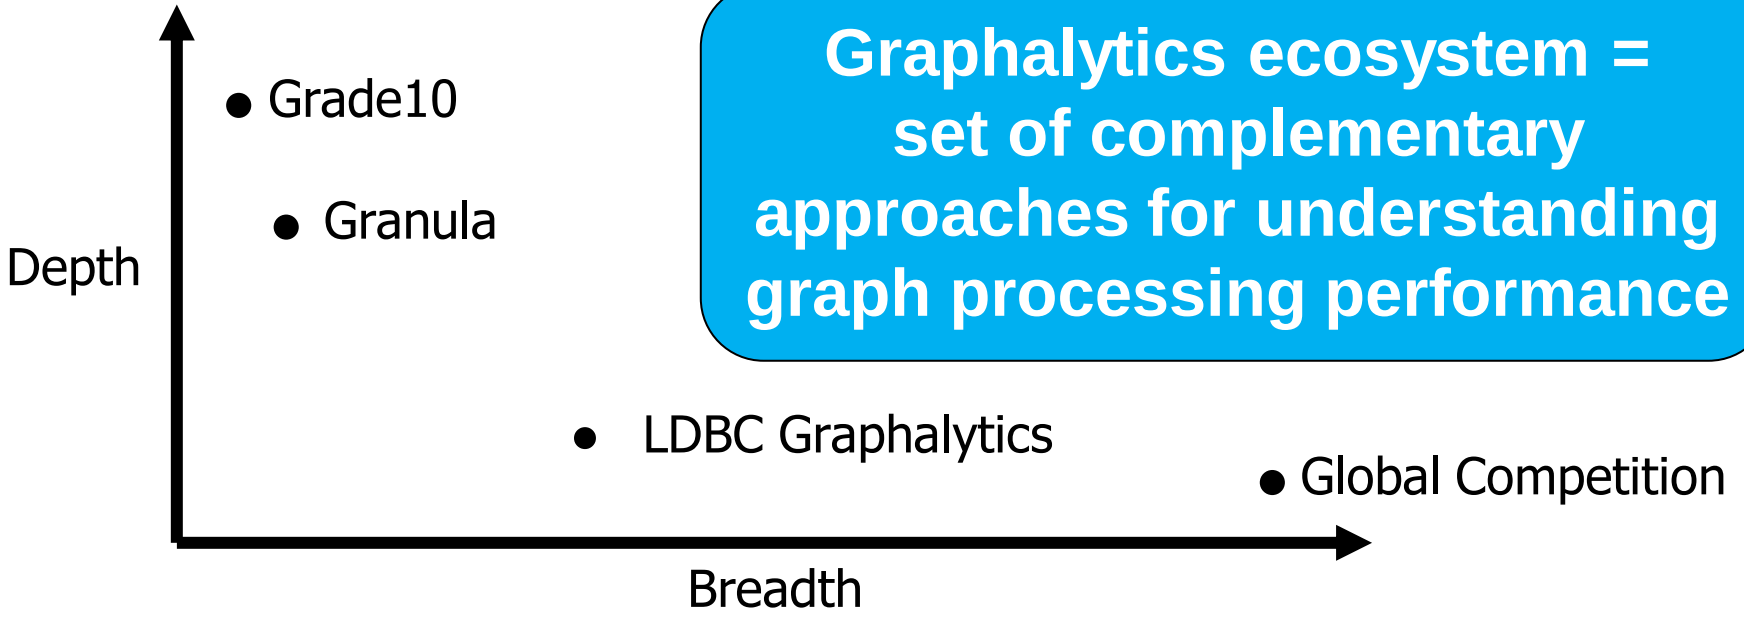

Graphalytics ecosystem = set of complementary approaches for understanding graph processing performance

The Graphalytics Ecosystem

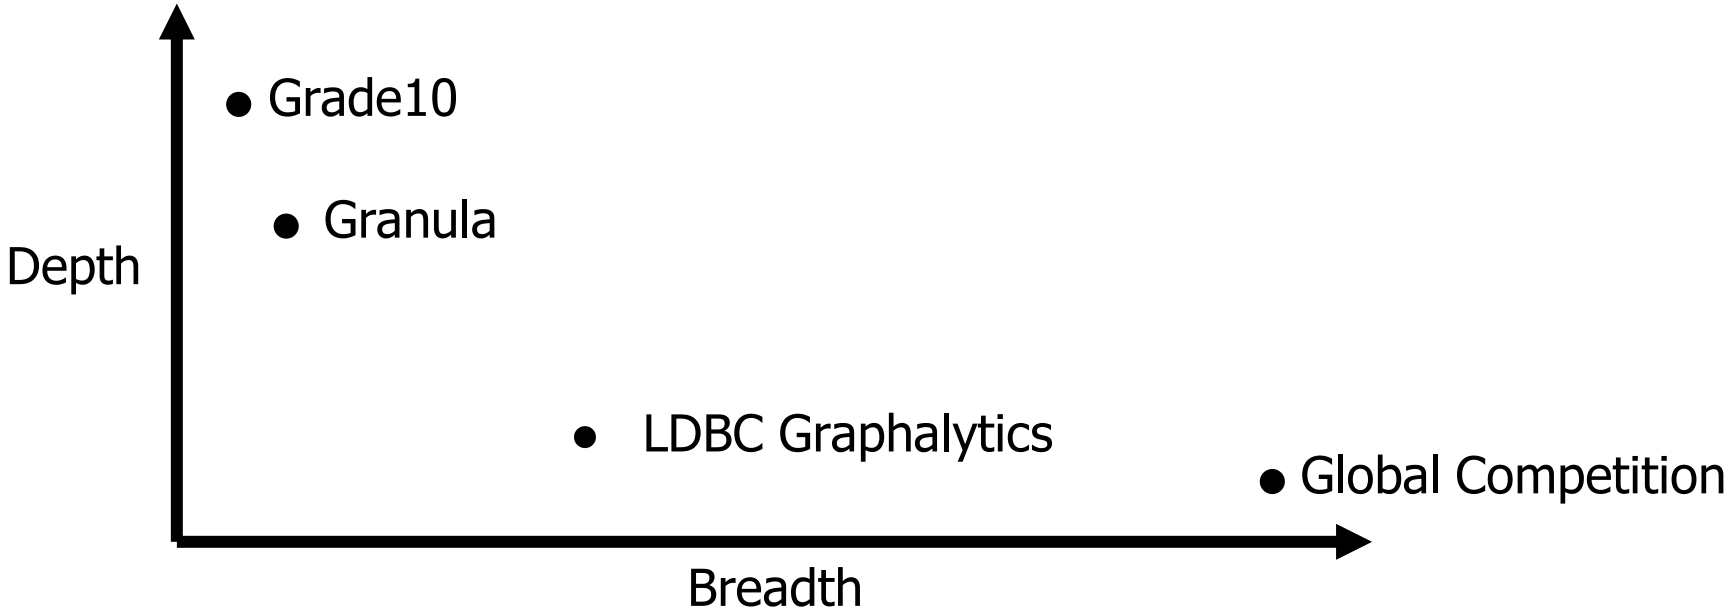

The Graphalytics Ecosystem

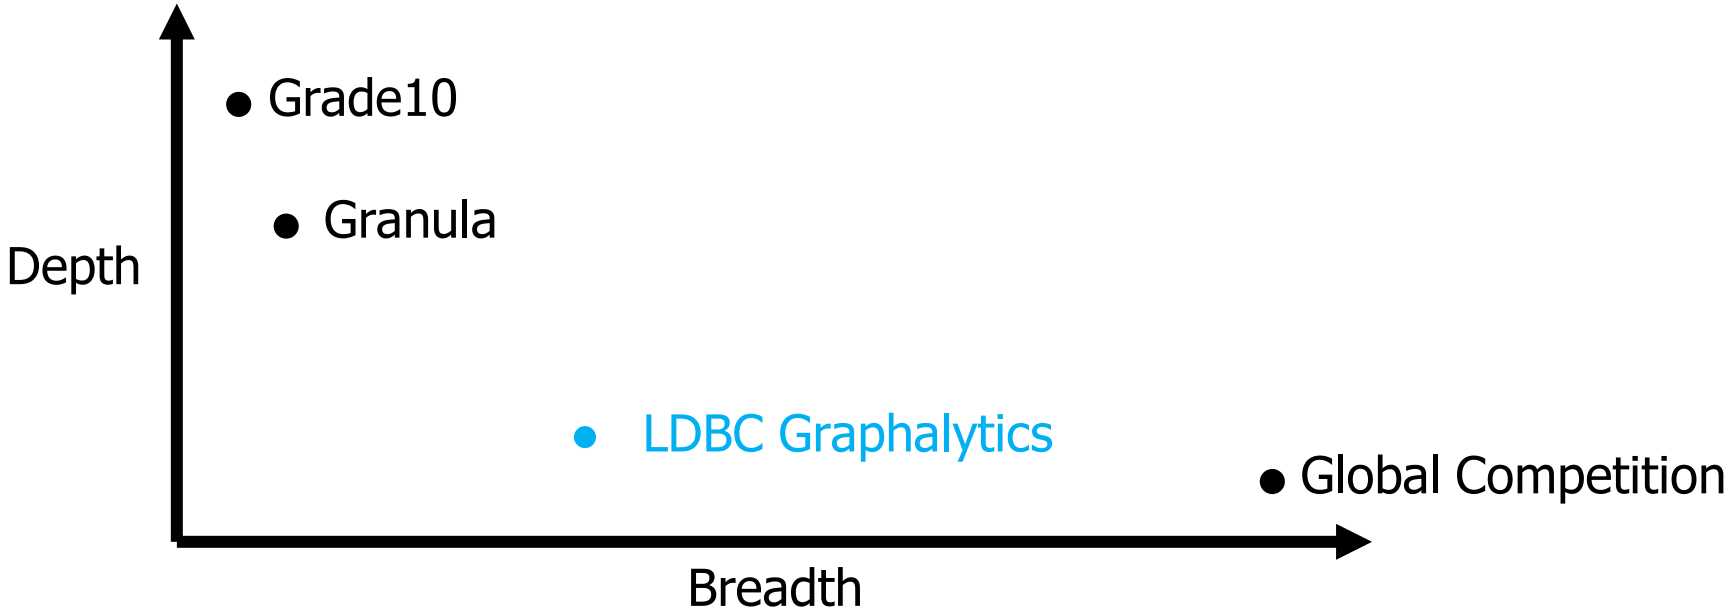

How to Compare the Performance of Graph Processing Platforms?

Typical approaches:

- Platform-centric comparative studies
  - Prove the superiority of a given system, limited set of metrics
- Benchmarks (Graph500, GreenGraph500, GraphBench, XGDBench, …)
  - Issues with representativeness, systems covered, metrics, …

How to Compare the Performance of Graph Processing Platforms?

LDBC Graphalytics = comprehensive benchmarking suite for graph processing across many platforms

A. Iosup et al. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. In *PVLDB,* vol. 9.13, 2016.

Graphalytics, in a Nutshell

- An LDBC benchmark
- Advanced benchmarking harness
- Many classes of algorithms used in practice
- Diverse real and synthetic datasets
- Diverse set of experiments representative for practice
- Renewal process to keep the workload relevant
- Enables comparison of many platforms, community-driven and industrial

https://graphalytics.org

Main Finding

Performance of graph processing is a non-trivial function of (Platform, Algorithm, Dataset, …), the PAD triangle

The Algorithm has a large impact

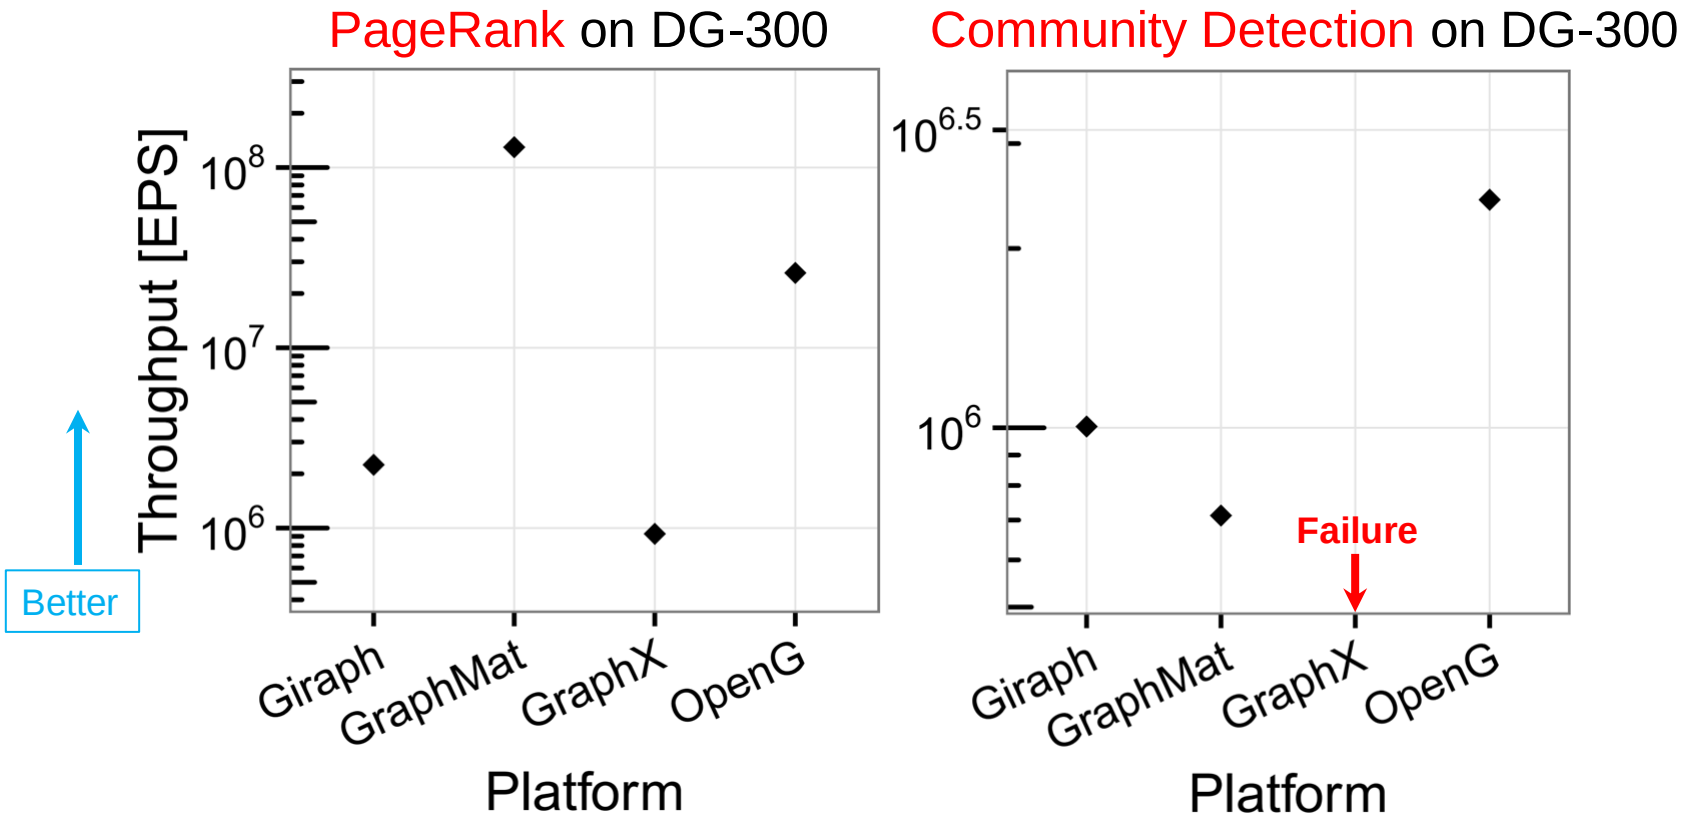

The Algorithm has a large impact

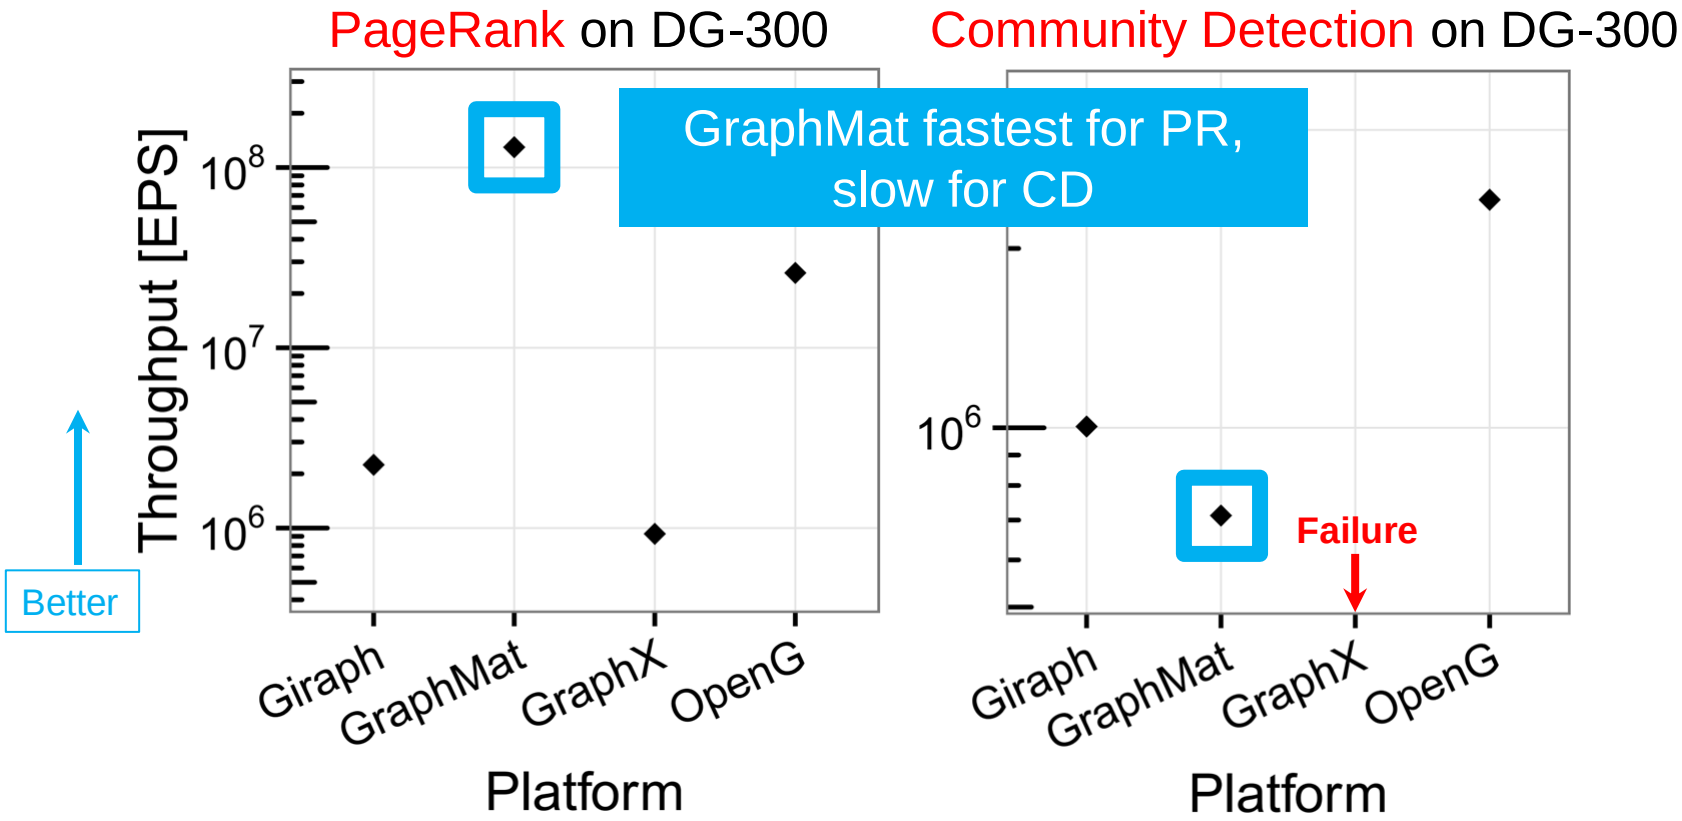

The Graphalytics Ecosystem

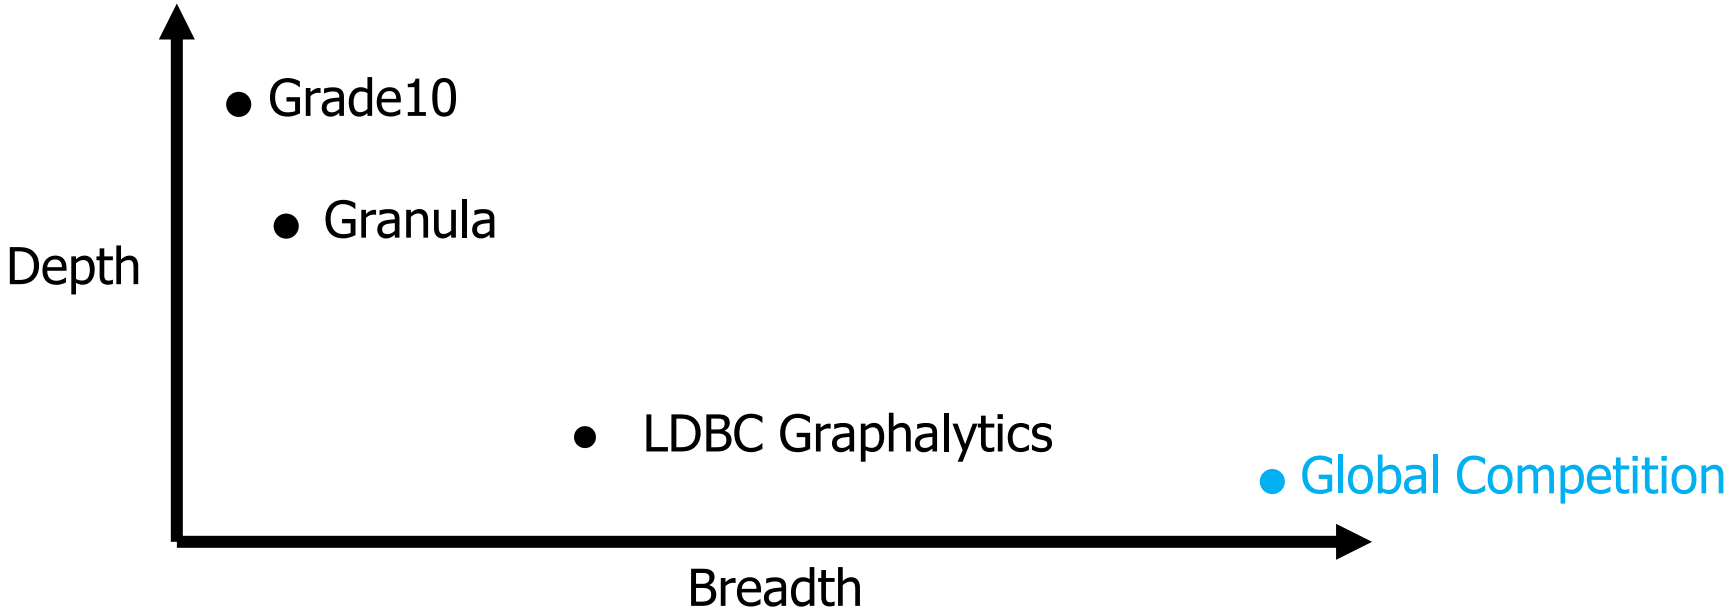

Graphalytics Global Competition

graphalytics.org/competition

2018 Spring competition #1

16 nodes, 32 threads - L graphs (public)

EDGES AND VERTICES PER SECOND

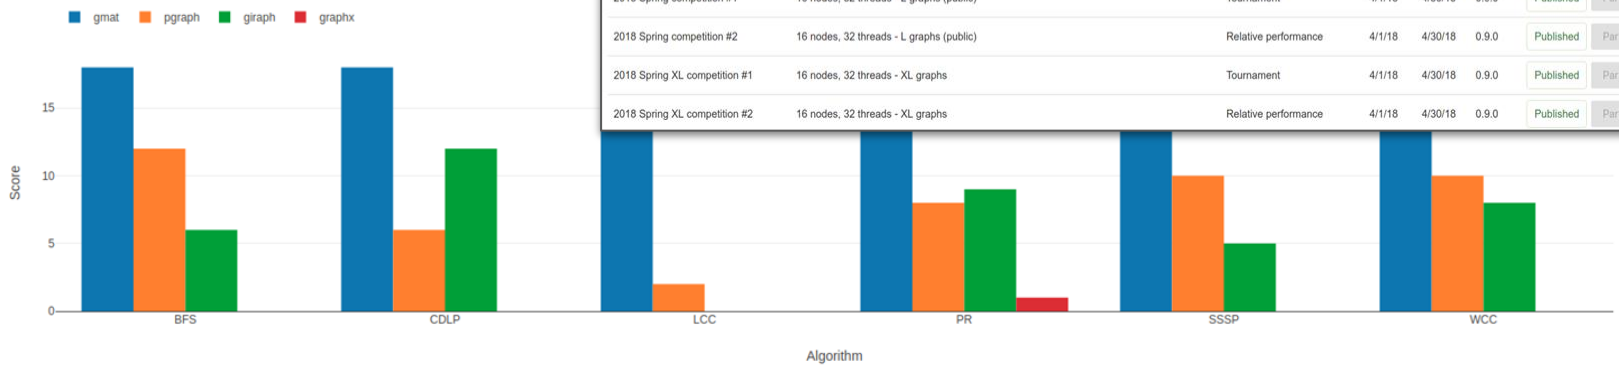

Competitions

At this moment, the Graphalytics Global Competition defines two types of competitions: (1) the *Tournament* competition and the (2) *relative-performance* competition.

A detailed description and explanation of both competitions can be found in the specification of different Graphalytics competitions report.

Name	Description	Type	Start	End	Spec	Status		
2018 Spring competition #1	16 nodes, 32 threads - L graphs (public)	Tournament	4/1/18	4/30/18	0.9.0	Published	Participate	Results
2018 Spring competition #2	16 nodes, 32 threads - L graphs (public)	Relative performance	4/1/18	4/30/18	0.9.0	Published	Participate	Results
2018 Spring XL competition #1	16 nodes, 32 threads - XL graphs	Tournament	4/1/18	4/30/18	0.9.0	Published	Participate	Results
2018 Spring XL competition #2	16 nodes, 32 threads - XL graphs	Relative performance	4/1/18	4/30/18	0.9.0	Published	Participate	Results

Rank	System name	Total score	BFS	CDLP	LCC	PR	SSSP	WCC
No. 1	GraphMat	105	18	18	18	18	15	18
No. 2	PowerGraph	48	12	6	2	8	10	10
No. 3	Giraph	40	6	12	0	9	5	8
No. 4	GraphX	1	0	0	0	1	0	0

Graphalytics Global Competition

graphalytics.org/competition

2018 Spring competition #1

16 nodes, 32 threads - L graphs (public)

EDGES AND VERTICES PER SECOND

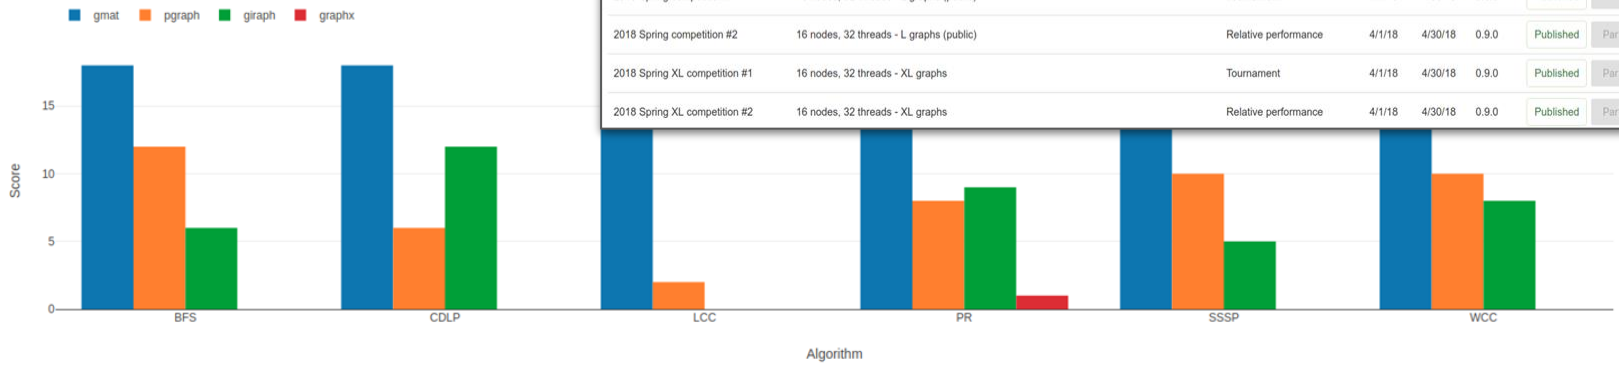

Competitions

At this moment, the Graphalytics Global Competition defines two types of competitions: (1) the *Tournament* competition and the (2) *relative-performance* competition.

A detailed description and explanation of both competitions can be found in the specification of different Graphalytics competitions report.

Name	Description	Type	Start	End	Spec	Status		
2018 Spring competition #1	16 nodes, 32 threads - L graphs (public)	Tournament	4/1/18	4/30/18	0.9.0	Published	Participate	Results
2018 Spring competition #2	16 nodes, 32 threads - L graphs (public)	Relative performance	4/1/18	4/30/18	0.9.0	Published	Participate	Results
2018 Spring XL competition #1	16 nodes, 32 threads - XL graphs	Tournament	4/1/18	4/30/18	0.9.0	Published	Participate	Results
2018 Spring XL competition #2	16 nodes, 32 threads - XL graphs	Relative performance	4/1/18	4/30/18	0.9.0	Published	Participate	Results

Rank	System name	Total score	BFS	CDLP	LCC	PR	SSSP	WCC
No. 1								18
No. 2								10
No. 3								8
No. 4	GraphX	1	0	0	0	1	0	0

Systematic and periodic comparison

Graphalytics Global Competition

Breadth-first search

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	2,185,887 kEVPS +3	2,170,844 kEVPS +3	438,309 kEVPS +3	1,930,948 kEVPS +3	461,637 kEVPS +3	2,549,718 kEVPS +3
PowerGraph	12	92,709 kEVPS +2	95,225 kEVPS +2	14,768 kEVPS +2	79,172 kEVPS +2	17,197 kEVPS +2	107,126 kEVPS +2
Giraph	6	35,876 kEVPS +1	38,133 kEVPS +1	8,455 kEVPS +1	38,291 kEVPS +1	9,853 kEVPS +1	46,299 kEVPS +1
GraphX	0	5,722 kEVPS +0	5,423 kEVPS +0	2,389 kEVPS +0	3,499 kEVPS +0	2,806 kEVPS +0	5,402 kEVPS +0

Community detection using label propagation

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	40,421 kEVPS +3	41,133 kEVPS +3	31,151 kEVPS +3	41,883 kEVPS +3	32,051 kEVPS +3	42,024 kEVPS +3
Giraph	12	8,950 kEVPS +2	9,104 kEVPS +2	5,151 kEVPS +2	8,179 kEVPS +2	5,414 kEVPS +2	8,714 kEVPS +2
PowerGraph	6	2,704 kEVPS +1	2,821 kEVPS +1	1,407 kEVPS +1	6,949 kEVPS +1	1,411 kEVPS +1	3,046 kEVPS +1
GraphX	0	- +0	- +0	- +0	- +0	- +0	- +0

Graphalytics Global Competition

Breadth-first search

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	2,185,887 kEVPS +3	2,170,844 kEVPS +3	438,309 kEVPS +3	1,930,948 kEVPS +3	461,637 kEVPS +3	2,549,718 kEVPS +3
PowerGraph	12	92,709 kEVPS +2	95,225 kEVPS +2	14,768 kEVPS +2	79,172 kEVPS +2	17,197 kEVPS +2	107,126 kEVPS +2
Giraph	6	35,876 kEVPS +1	38,133 kEVPS +1	8,455 kEVPS +1	38,291 kEVPS +1	9,853 kEVPS +1	46,299 kEVPS +1
GraphX	0	5,722 kEVPS +0	5,423 kEVPS +0	2,389 kEVPS +0	3,499 kEVPS +0	2,806 kEVPS +0	5,402 kEVPS +0

Community detection using label propagation

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	40,421 kEVPS +3	41,133 kEVPS +3	31,151 kEVPS +3	41,883 kEVPS +3	32,051 kEVPS +3	42,024 kEVPS +3
Giraph	12	8,950 kEVPS +2	9,104 kEVPS +2	5,151 kEVPS +2	8,179 kEVPS +2	5,414 kEVPS +2	8,714 kEVPS +2
PowerGraph	6	2,704 kEVPS +1	2,821 kEVPS +1	1,407 kEVPS +1	6,949 kEVPS +1	1,411 kEVPS +1	3,046 kEVPS +1
GraphX	0	- +0	- +0	- +0	- +0	- +0	- +0

Graphalytics Global Competition

Breadth-first search

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	2,185,887 kEVPS +3	2,170,844 kEVPS +3	438,309 kEVPS +3	1,930,948 kEVPS +3	461,637 kEVPS +3	2,549,718 kEVPS +3
PowerGraph	12	92,709 kEVPS +2	95,225 kEVPS +2	14,768 kEVPS +2	79,172 kEVPS +2	17,197 kEVPS +2	107,126 kEVPS +2
Giraph	6	35,876 kEVPS +1	38,133 kEVPS +1	8,455 kEVPS +1	38,291 kEVPS +1	9,853 kEVPS +1	46,299 kEVPS +1
GraphX	0	5,722 kEVPS +0	5,423 kEVPS +0	2,389 kEVPS +0	3,499 kEVPS +0	2,806 kEVPS +0	5,402 kEVPS +0

Community detection using label propagation

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	40,421 kEVPS +3	41,133 kEVPS +3	31,151 kEVPS +3	41,883 kEVPS +3	32,051 kEVPS +3	42,024 kEVPS +3
Giraph	12	8,950 kEVPS +2	9,104 kEVPS +2	5,151 kEVPS +2	8,179 kEVPS +2	5,414 kEVPS +2	8,714 kEVPS +2
PowerGraph	6	2,704 kEVPS +1	2,821 kEVPS +1	1,407 kEVPS +1	6,949 kEVPS +1	1,411 kEVPS +1	3,046 kEVPS +1
GraphX	0	- +0	- +0	- +0	- +0	- +0	- +0

Graphalytics Global Competition

Breadth-first search

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	2,185,887 kEVPS +3	[illegible] +3	438,309 kEVPS +3	1,930,948 kEVPS +3	461,637 kEVPS +3	2,549,718 kEVPS +3
PowerGraph	12	92,709 kEVPS +2	95,225 kEVPS +2	14,768 kEVPS +2	79,172 kEVPS +2	17,197 kEVPS +2	107,126 kEVPS +2
Giraph	6	35,876 kEVPS +1	38,133 kEVPS +1	8,455 kEVPS +1	38,291 kEVPS +1	9,853 kEVPS +1	46,299 kEVPS +1
GraphX	0	5,722 kEVPS +0	5,423 kEVPS +0	2,389 kEVPS +0	3,499 kEVPS +0	2,806 kEVPS +0	5,402 kEVPS +0

Community detection using label propagation

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	40,421 kEVPS +3	41,133 kEVPS +3	31,151 kEVPS +3	41,883 kEVPS +3	32,051 kEVPS +3	42,024 kEVPS +3
Giraph	12	8,950 kEVPS +2	9,104 kEVPS +2	5,151 kEVPS +2	8,179 kEVPS +2	5,414 kEVPS +2	8,714 kEVPS +2
PowerGraph	6	2,704 kEVPS +1	2,821 kEVPS +1	1,407 kEVPS +1	6,949 kEVPS +1	1,411 kEVPS +1	3,046 kEVPS +1
GraphX	0	- +0	- +0	- +0	- +0	- +0	- +0

Graphalytics Global Competition

Breadth-first search

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	2,185,887 kEVPS +3	2,170,844 kEVPS +3	438,309 kEVPS +3	1,930,948 kEVPS +3	461,637 kEVPS +3	2,549,718 kEVPS +3
PowerGraph	12	92,709 kEVPS +2	95,225 kEVPS +2	14,768 kEVPS +2	79,172 kEVPS +2	17,197 kEVPS +2	107,126 kEVPS +2
Giraph	6	35,876 kEVPS +1	38,133 kEVPS +1	8,455 kEVPS +1	38,291 kEVPS +1	9,853 kEVPS +1	46,299 kEVPS +1
GraphX	0	5,722 kEVPS +0	5,423 kEVPS +0	2,389 kEVPS +0	3,499 kEVPS +0	2,806 kEVPS +0	5,402 kEVPS +0

Community detection using label propagation

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	40,421 kEVPS +3	41,133 kEVPS +3	31,151 kEVPS +3	41,883 kEVPS +3	32,051 kEVPS +3	42,024 kEVPS +3
Giraph	12	8,950 kEVPS +2	9,104 kEVPS +2	5,151 kEVPS +2	8,179 kEVPS +2	5,414 kEVPS +2	8,714 kEVPS +2
PowerGraph	6	2,704 kEVPS +1	2,821 kEVPS +1	1,407 kEVPS +1	6,949 kEVPS +1	1,411 kEVPS +1	3,046 kEVPS +1
GraphX	0	- +0	- +0	- +0	- +0	- +0	- +0

Graphalytics Global Competition

Breadth-first search

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	2,185,887 kEVPS +3	2,170,844 kEVPS +3	438,309 kEVPS +3	1,930,948 kEVPS +3	461,637 kEVPS +3	2,549,718 kEVPS +3
PowerGraph	12	92,709 kEVPS +2	95,225 kEVPS +2	14,768 kEVPS +2	79,172 kEVPS +2	17,197 kEVPS +2	107,126 kEVPS +2
Giraph	6	35,876 kEVPS +1	38,133 kEVPS +1	8,455 kEVPS +1	38,291 kEVPS +1	9,853 kEVPS +1	46,299 kEVPS +1
GraphX	0	5,722 kEVPS +0	5,423 kEVPS +0	2,389 kEVPS +0	3,499 kEVPS +0	2,806 kEVPS +0	5,402 kEVPS +0

Community detection using label propagation

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	40,421 kEVPS +3	41,133 kEVPS +3	31,151 kEVPS +3	41,883 kEVPS +3	32,051 kEVPS +3	42,024 kEVPS +3
Giraph	12	8,950 kEVPS +2	9,104 kEVPS +2	5,151 kEVPS +2	8,179 kEVPS +2	5,414 kEVPS +2	8,714 kEVPS +2
PowerGraph	6	2,704 kEVPS +1	2,821 kEVPS +1	1,407 kEVPS +1	6,949 kEVPS +1	1,411 kEVPS +1	3,046 kEVPS +1
GraphX	0	- +0	- +0	- +0	- +0	- +0	- +0

Graphalytics Global Competition

Breadth-first search

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	2,185,887 kEVPS +3	2,170,844 kEVPS +3	438,309 kEVPS +3	1,930,948 kEVPS +3	461,637 kEVPS +3	2,549,718 kEVPS +3
PowerGraph	12	92,709 kEVPS +2	95,225 kEVPS +2	14,768 kEVPS +2	79,172 kEVPS +2	17,197 kEVPS +2	107,126 kEVPS +2
Giraph	6	35,876 kEVPS +1	38,133 kEVPS +1	8,455 kEVPS +1	38,291 kEVPS +1	9,853 kEVPS +1	46,299 kEVPS +1
GraphX	0	5,722 kEVPS +0	5,423 kEVPS +0	2,389 kEVPS +0	3,499 kEVPS +0	2,806 kEVPS +0	5,402 kEVPS +0

Community detection using label propagation

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	40,421 kEVPS +3	41,133 kEVPS +3	31,151 kEVPS +3	41,883 kEVPS +3	32,051 kEVPS +3	42,024 kEVPS +3
Giraph	12	8,950 kEVPS +2	9,104 kEVPS +2	5,151 kEVPS +2	8,179 kEVPS +2	5,414 kEVPS +2	8,714 kEVPS +2
PowerGraph	6	2,704 kEVPS +1	2,821 kEVPS +1	1,407 kEVPS +1	6,949 kEVPS +1	1,411 kEVPS +1	3,046 kEVPS +1
GraphX	0	- +0	- +0	- +0	- +0	- +0	- +0

Graphalytics Global Competition

Breadth-first search

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	2,185,887 kEVPS +3	2,170,844 kEVPS +3	438,309 kEVPS +3	1,930,948 kEVPS +3	461,637 kEVPS +3	2,549,718 kEVPS +3
PowerGraph	12	92,709 kEVPS +2	95,225 kEVPS +2	14,768 kEVPS +2	79,172 kEVPS +2	17,197 kEVPS +2	107,126 kEVPS +2
Giraph	6	35,876 kEVPS +1	38,133 kEVPS +1	8,455 kEVPS +1	38,291 kEVPS +1	9,853 kEVPS +1	46,299 kEVPS +1
GraphX	0	5,722 kEVPS +0	5,423 kEVPS +0	2,389 kEVPS +0	3,499 kEVPS +0	2,806 kEVPS +0	5,402 kEVPS +0

Community detection using label propagation

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	40,421 kEVPS +3	41,133 kEVPS +3	31,151 kEVPS +3	41,883 kEVPS +3	32,051 kEVPS +3	42,024 kEVPS +3
Giraph	12	8,950 kEVPS +2	9,104 kEVPS +2	5,151 kEVPS +2	8,179 kEVPS +2	5,414 kEVPS +2	8,714 kEVPS +2
PowerGraph	6	2,704 kEVPS +1	2,821 kEVPS +1	1,407 kEVPS +1	6,949 kEVPS +1	1,411 kEVPS +1	3,046 kEVPS +1
GraphX	0	- +0	- +0	- +0	- +0	- +0	- +0

Graphalytics Global Competition

Breadth-first search

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	2,185,887 kEVPS +3	2,170,844 kEVPS +3	438,309 kEVPS +3	1,930,948 kEVPS +3	461,637 kEVPS +3	2,549,718 kEVPS +3
PowerGraph	12	92,709 kEVPS +2	95,225 kEVPS +2	14,768 kEVPS +2	79,172 kEVPS +2	17,197 kEVPS +2	107,126 kEVPS +2
Giraph	6	35,876 kEVPS +1	38,133 kEVPS +1	8,455 kEVPS +1	38,291 kEVPS +1	9,853 kEVPS +1	46,299 kEVPS +1
GraphX	0	5,722 kEVPS +0	5,423 kEVPS +0	2,389 kEVPS +0	3,499 kEVPS +0	2,806 kEVPS +0	5,402 kEVPS +0

Community detection using label propagation

System name	Total score (EVPS)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
GraphMat	18	40,421 kEVPS +3	41,133 kEVPS +3	31,151 kEVPS +3	41,883 kEVPS +3	32,051 kEVPS +3	42,024 kEVPS +3
Giraph	12	8,950 kEVPS +2	9,104 kEVPS +2	5,151 kEVPS +2	8,179 kEVPS +2	5,414 kEVPS +2	8,714 kEVPS +2
PowerGraph	6	2,704 kEVPS +1	2,821 kEVPS +1	1,407 kEVPS +1	6,949 kEVPS +1	1,411 kEVPS +1	3,046 kEVPS +1
GraphX	0	- +0	- +0	- +0	- +0	- +0	- +0

The Graphalytics Ecosystem

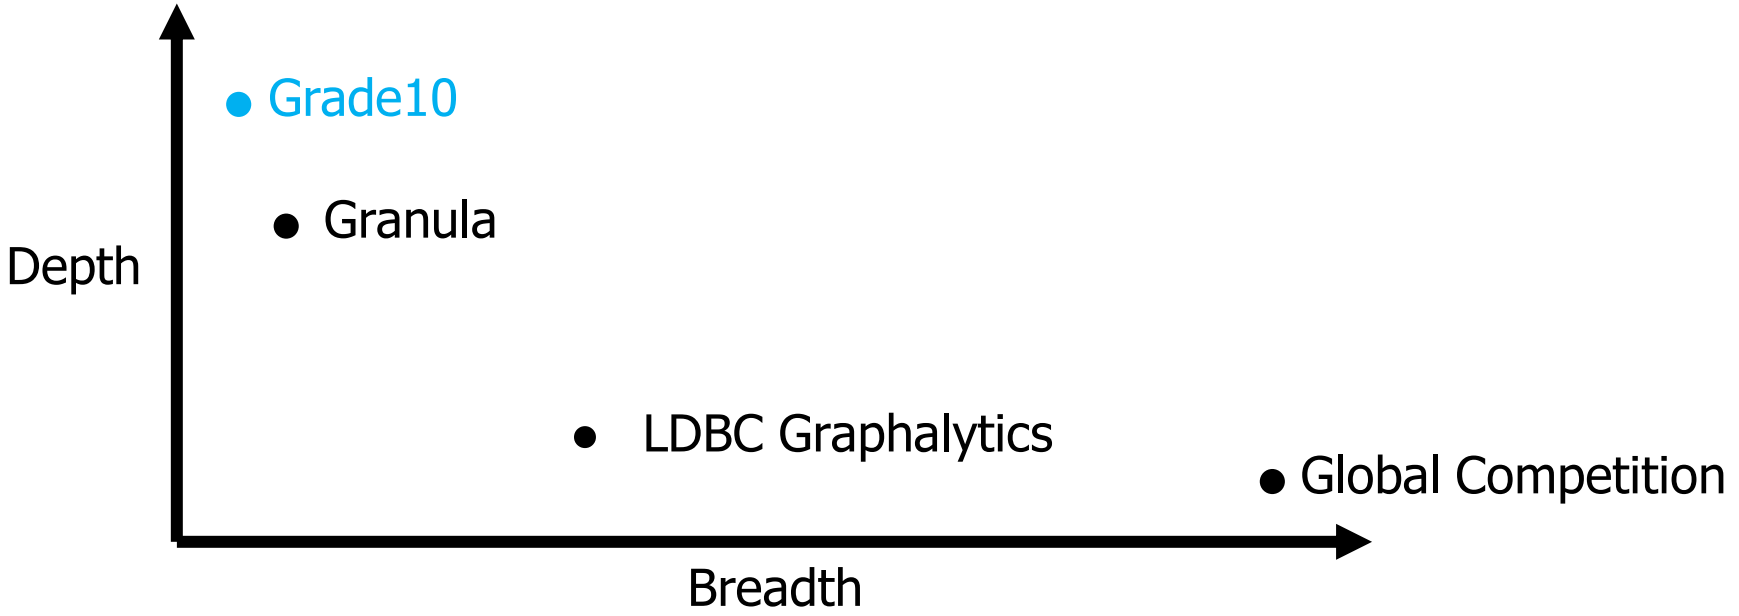

Performance Analysis

GOAL:

Identify bottlenecks and performance issues

Performance Analysis Techniques

Analytical modeling

$$calcNumSpillsFirstPass(N,F) = \begin{cases} N & , \text{if } N \leq F \\ F & , \text{if } (N-1)\ MOD\ (F-1) = 0 \\ (N-1)\ MOD\ (F-1)+1 & , \text{otherwise} \end{cases}$$

$$calcNumSpillsIntermMerge(N,F) = \begin{cases} 0 & , \text{if } N \leq F \\ P + \lfloor \frac{N-P}{F} \rfloor * F & , \text{if } N \leq F^2 \end{cases}$$
$$, where\ P = calcNumSpillsFirstPass(N,F)$$

Profiling

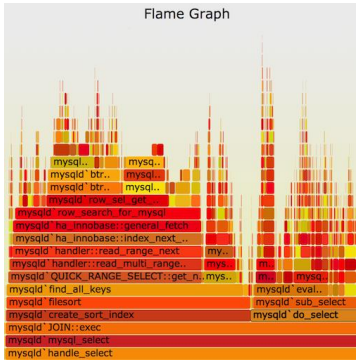

Tracing

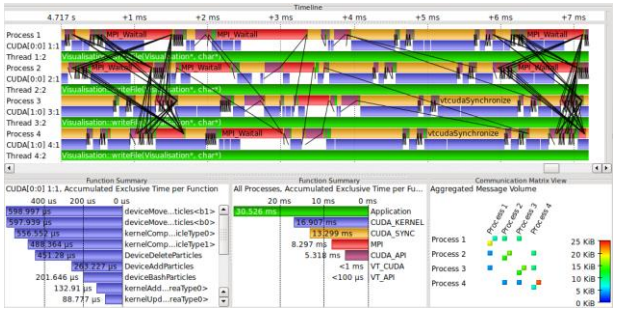

Performance Analysis Techniques

Analytical modeling

Profiling

Tracing

Infeasible for complex systems

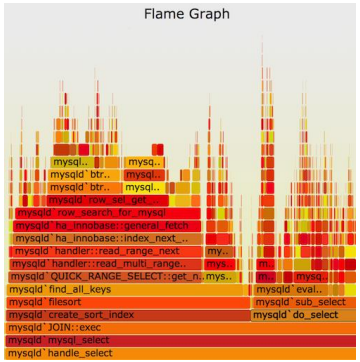

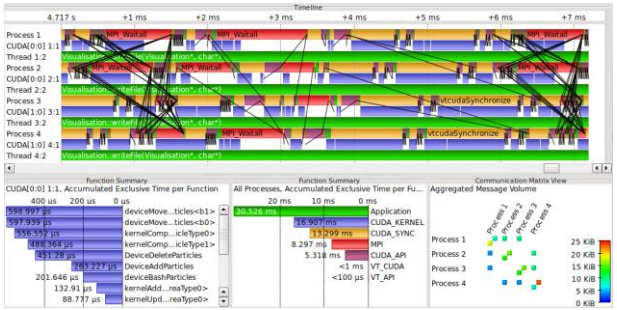

Performance Analysis Techniques

Analytical modeling

Profiling

Tracing

Infeasible for complex systems

Works, widely used, pinpoints issues

Performance Analysis Techniques

Analytical modeling

Profiling

Tracing

Infeasible for complex systems

Works, widely used, pinpoints issues

Expertise-driven, coupling with source code

Anatomy of an Application

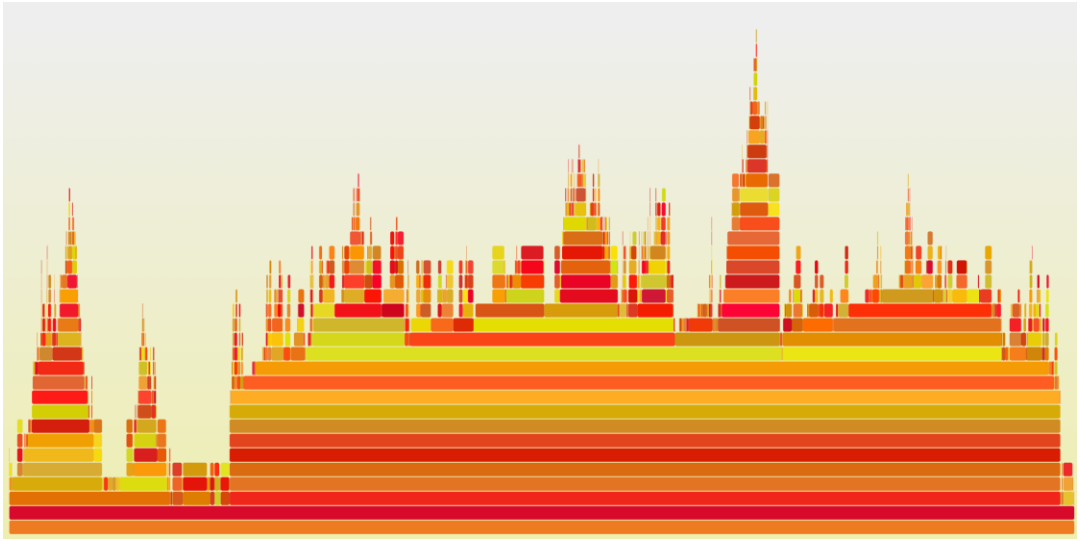

Anatomy of an Application

HPC

MPI, libraries

Our code

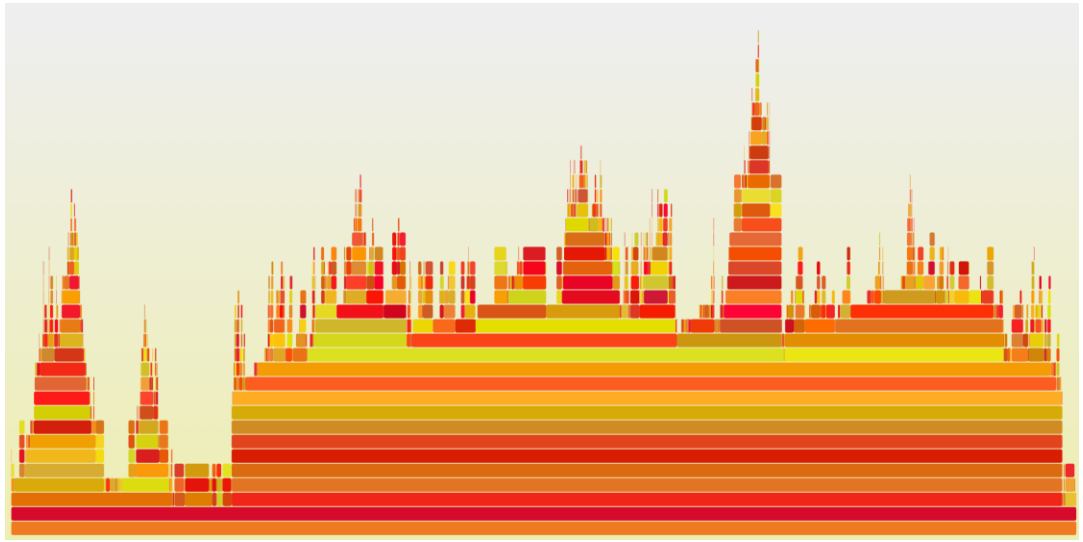

(Not to scale)

Anatomy of an Application

HPC

MPI, libraries

Our code

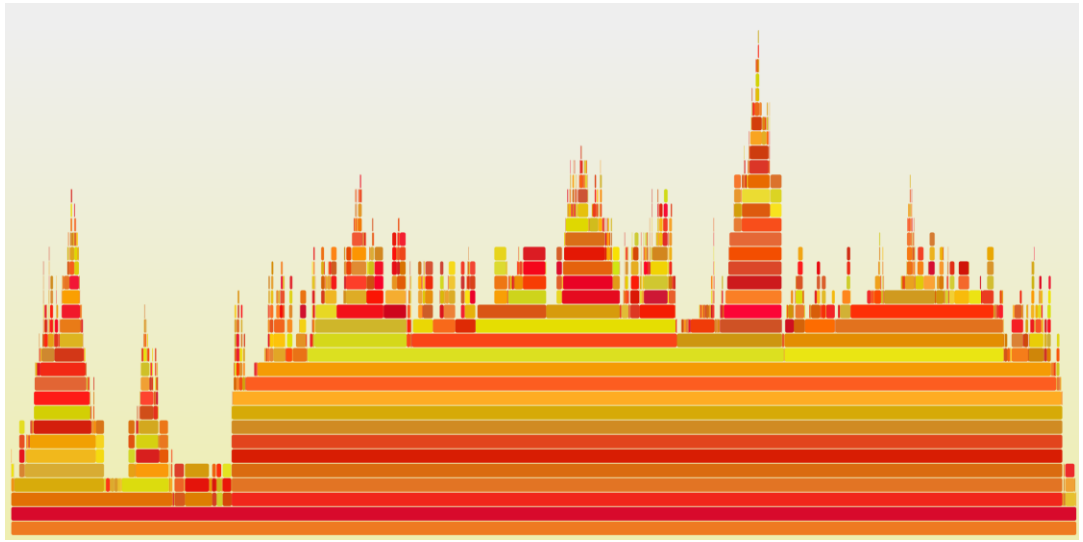

Graph Processing

Our code

Framework

(Not to scale)

Anatomy of a Graph Processing Application

Execution model

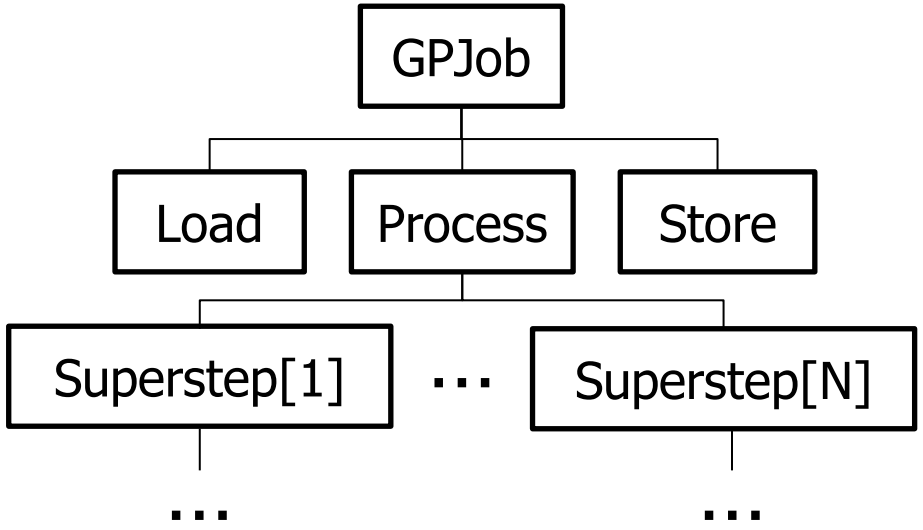

Instead of source code, map performance metrics to high-level (conceptual) stages of execution

Break down into substages to increase model accuracy and enable fine-grained performance analysis

Grade10: Automated Bottleneck Detection and Performance Issue Identification

Core idea:

Execution model of graph processing platform + resource monitoring data
=
Bottlenecks & performance issues of variable granularity

Grade10: Automated Bottleneck Detection and Performance Issue Identification

System under test

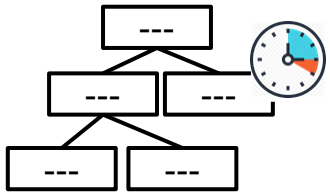

Execution model + Event logging

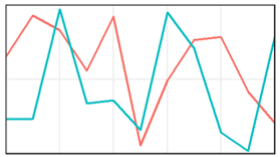

Monitoring

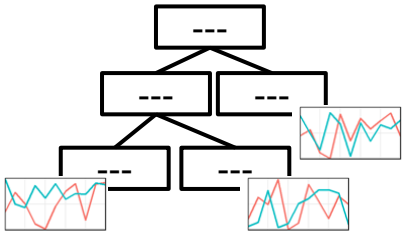

Resource attribution

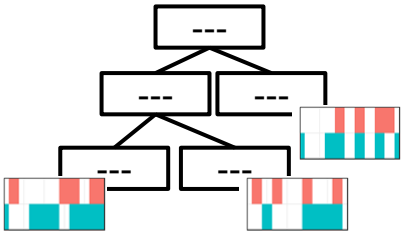

Bottleneck detection

Top bottlenecks:

Performance issue identification

Resource Attribution

Given a resource's observed usage and a set of (parallel) phases,
how much of the resource did each phase use?

Resource Attribution

Given a resource's observed usage and a set of (parallel) phases,
how much of the resource did each phase use?

Short answer: each phase used an equal share, unless we are told
how much the phase is expected to use.

Example Result: Analysing a Giraph Job

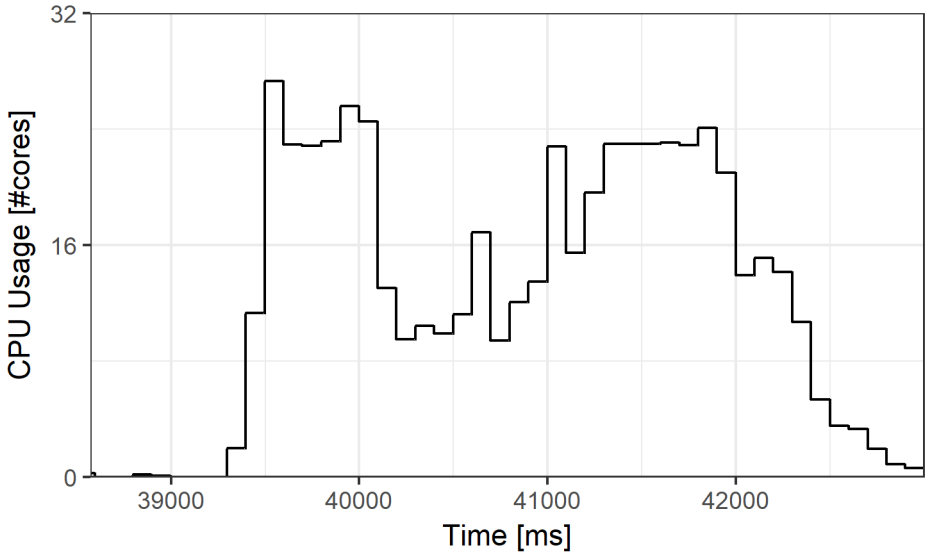

Example Result: Analysing a Giraph Job

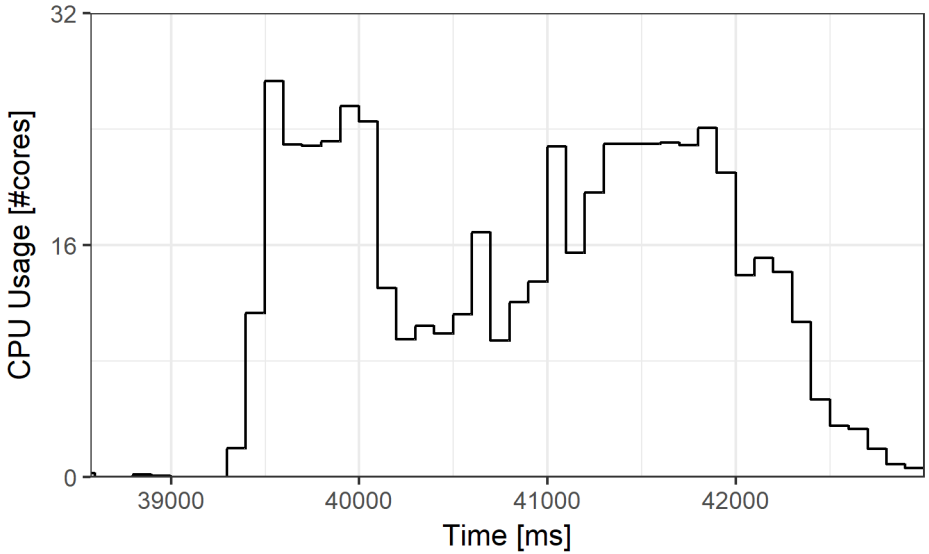

Initial observation:
CPU usage < 32 cores (= 100%),
so no bottleneck?

Example Result: Mapping to Execution Model

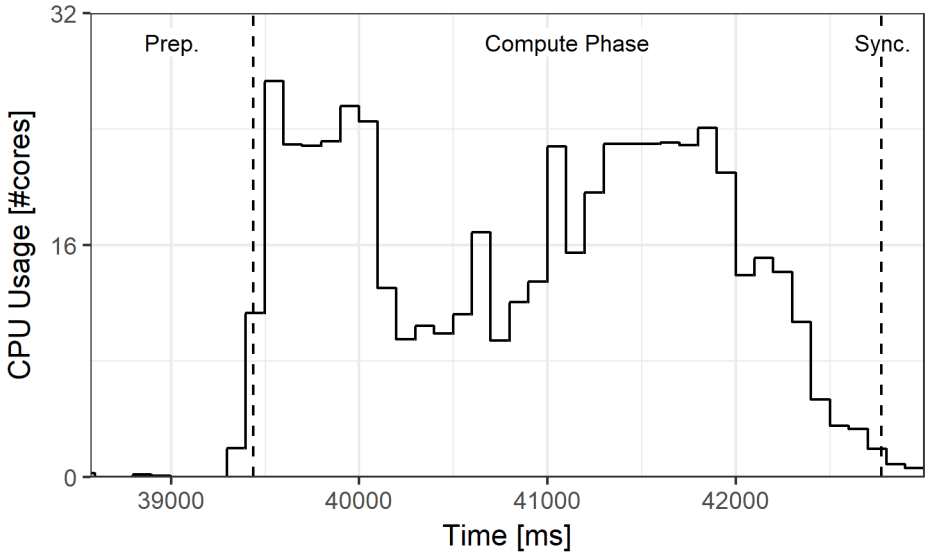

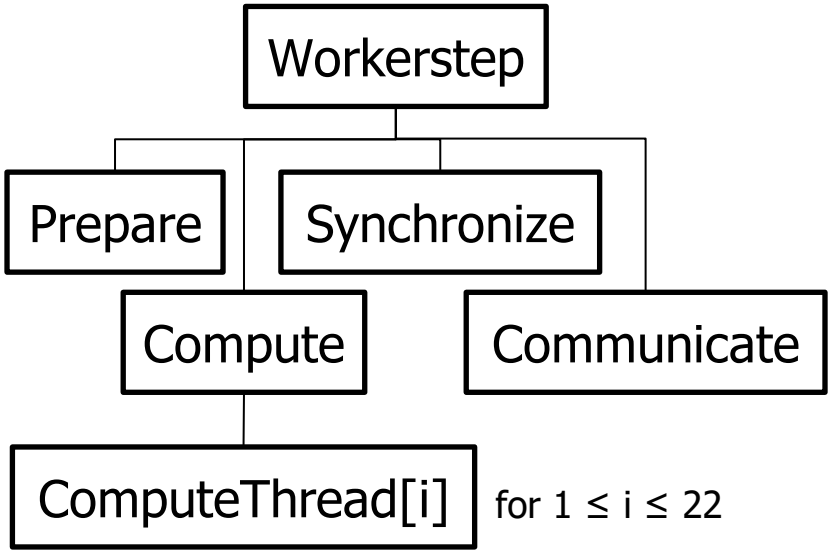

Example Result: Resource Attribution

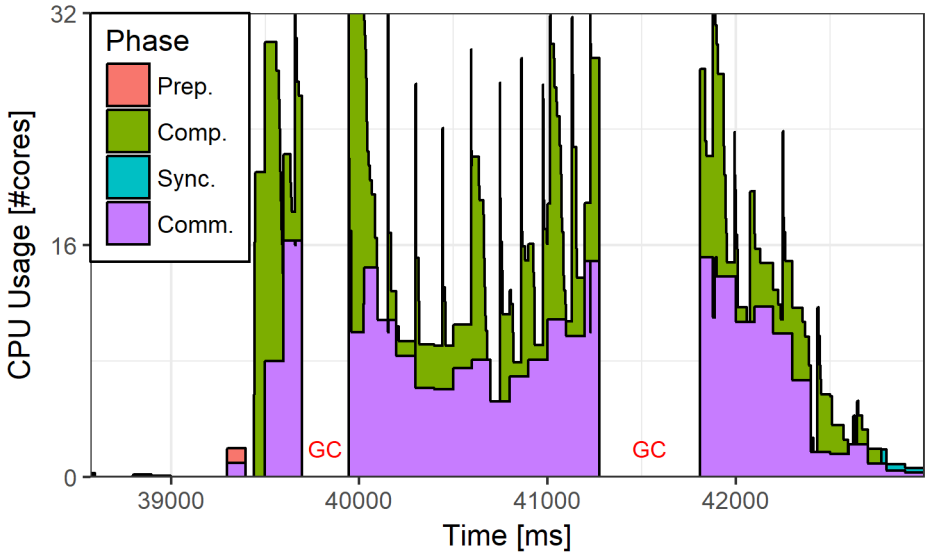

Observation:
CPU usage for Compute is bursty; ComputeThreads do not use CPU when waiting on full message queue/garbage collection

Example Result: Bottleneck Detection

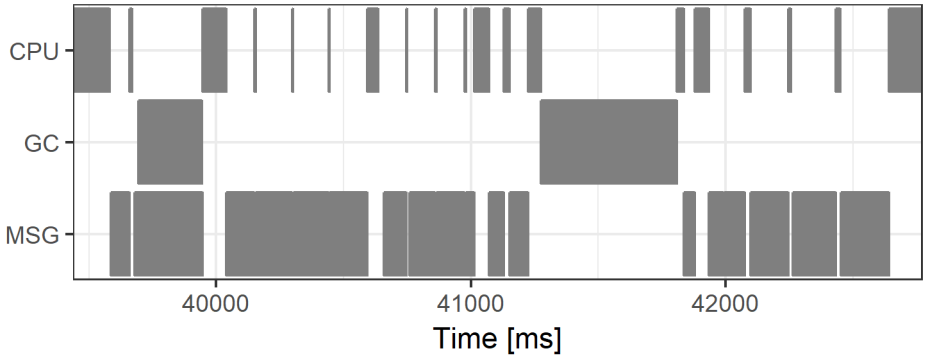

Bottleneck profile of an average ComputeThread

Example Result: Performance Issues

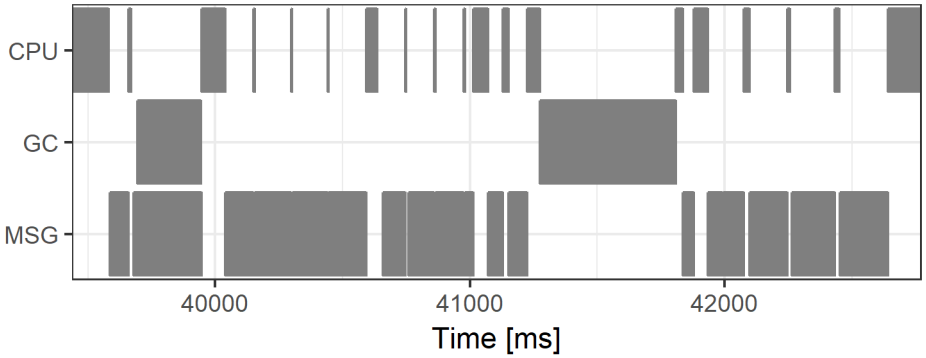

Top bottlenecks for Compute:
- Message queue full: 1768 ms
- Garbage collect: 781 ms
- CPU: 748 ms
- None: 0 ms

Focus on reducing:
- Communication overhead
- GC overhead (good luck!)

Take-home Message

The Graphalytics ecosystem provides breadth and depth in understanding graph processing performance.

LDBC Graphalytics: comprehensive benchmarking suite for graph processing across many platforms.

Global Competition: a systematic and periodic comparison of graph processing platforms.

Grade10: automated bottleneck detection and performance issue identification.

https://graphalytics.org

The Graphalytics Ecosystem

From Competitions to Performance Analysis

June 20, 2018


Ahmed Musaafir


Tim Hegeman


Wing Lung Ngai


Dr. Alexandru Uta


Prof.dr.ir. Alexandru Iosup

VU VRIJE UNIVERSITEIT AMSTERDAM

TU Delft Delft University of Technology

atlarge-research.com
graphalytics.org

Further Reading

A. Iosup et al. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. In *PVLDB,* vol. 9.13, 2016.

W.L. Ngai et al. Granula: Toward Fine-grained Performance Analysis of Large-scale Graph Processing Platforms. GRADES@SIGMOD/PODS 2017: 8:1-8:6

Graphalytics Global Competition (results & specification) is accessible via https://graphalytics.org