VU VRIJE UNIVERSITEIT AMSTERDAM

NWO

LDBC **Graphalytics**

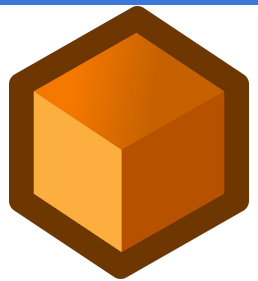

12th LDBC TUC Meeting (July 5th, 2019)

Graphalytics team, including Ahmed Musaafir, Tim Hegeman, Alexandru Uta, Alexandru Iosup

AtLarge Massivizing Computer Systems research group, Vrije Universiteit Amsterdam, the Netherlands

https://graphalytics.org - https://atlarge-research.com

Graphs, Many Graph-Processing Platforms

Google

Linked in

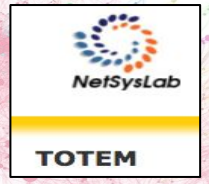

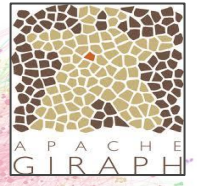

Intel Graphmat

nvGRAPH

Analyse large graphs

Which platforms perform well?

Why? How to improve?

The Data Deluge: Large-scale Graphs

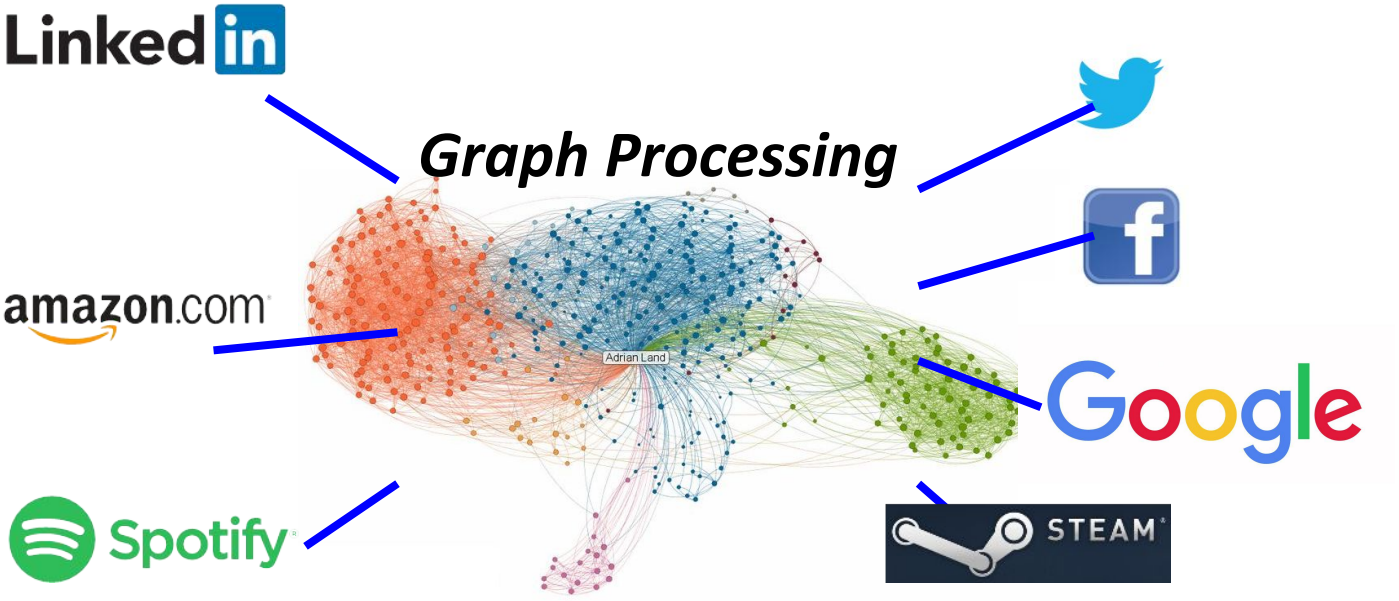

The Data Deluge: Large-scale Graphs

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

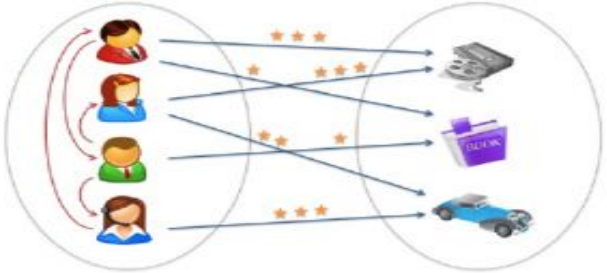

Navigation systems

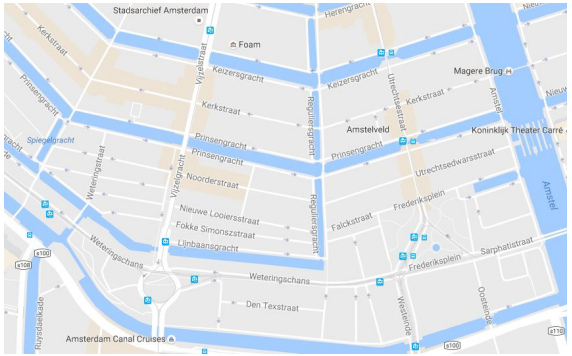

The Data Deluge: Large-scale Graphs

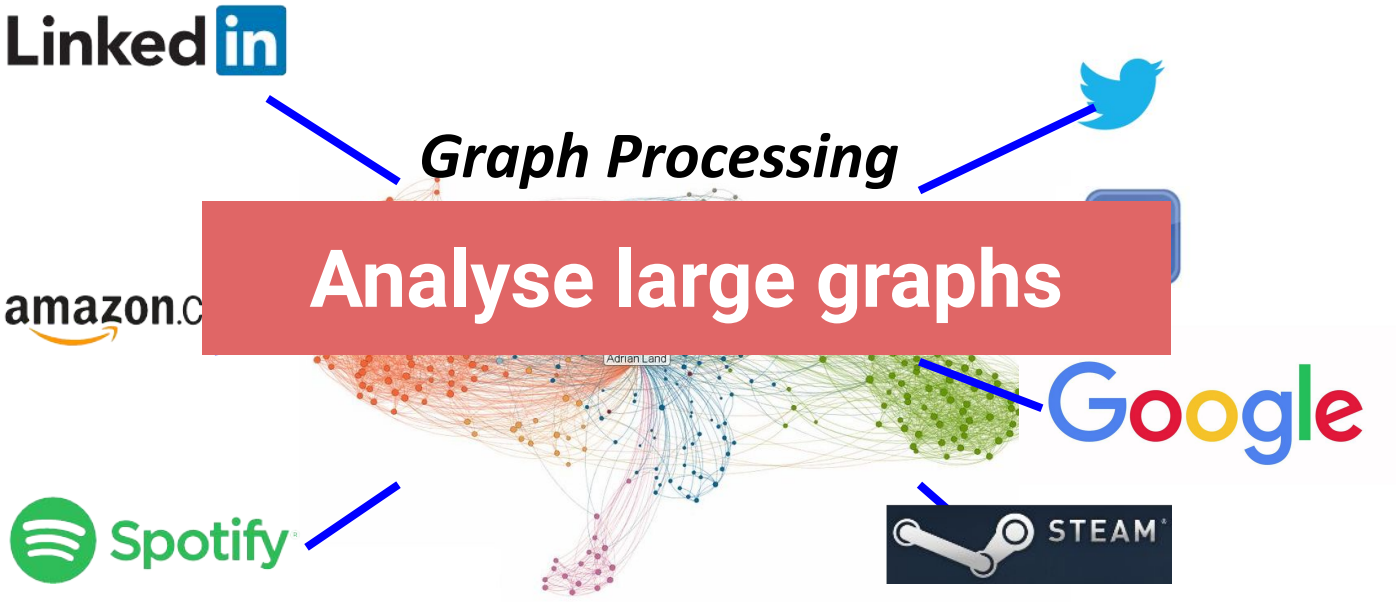

Graph Processing Platforms

Intel Graphmat

Gemini

IBM System G

Gunrock

nvGRAPH

JoyGraph

HyGraph

Graph Processing Platforms

Intel Graphmat

Gemini

IBM System G

Gunrock

nvGRAPH

JoyGraph

HyGraph

Which platforms perform well?

Graph Processing Platforms

Intel Graphmat

Gemini

IBM System G

Gunrock

nvGRAPH

JoyGraph

HyGraph

Which platforms perform well?

Why? How can they be improved?

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 components for understanding graph processing performance.

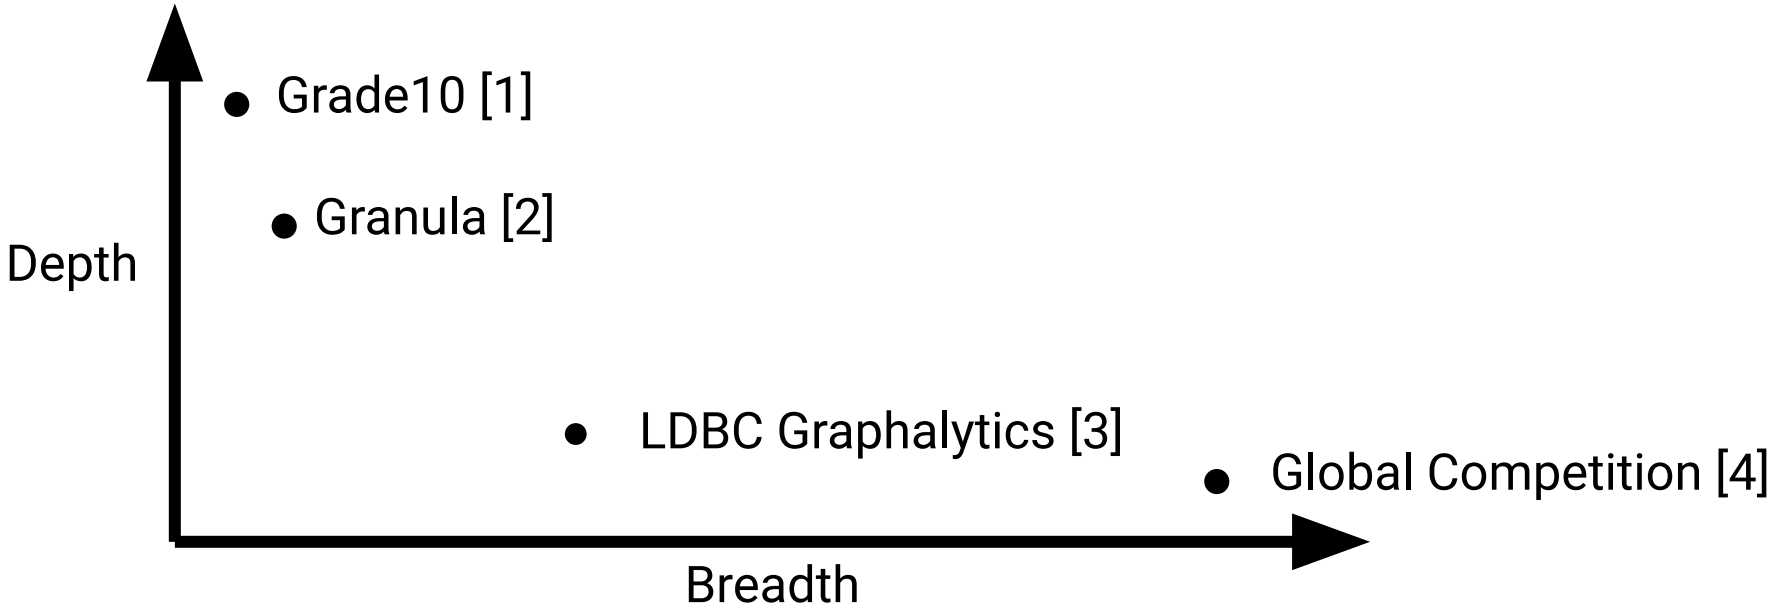

[1] Hegeman, Iosup. Experimental Performance Analysis of Graph Analytics Frameworks. MA thesis. TU Delft, 2018.
[2] Ngai et al., Toward Fine-grained Performance Analysis of Large-scale Graph Processing Platforms. GRADES@SIGMOD/PODS 2017: 8:1-6
[3] Iosup, et al.,. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. PVLDB 9(13): 1317-1328, 2016)
[4] Specification of different Graphalytics Competitions https://graphalytics.org/assets/spec-graphalytics-competitions.pdf, 2018

The Graphalytics Ecosystem

Graphalytics ecosystem: set of complementary components for understanding graph processing performance.

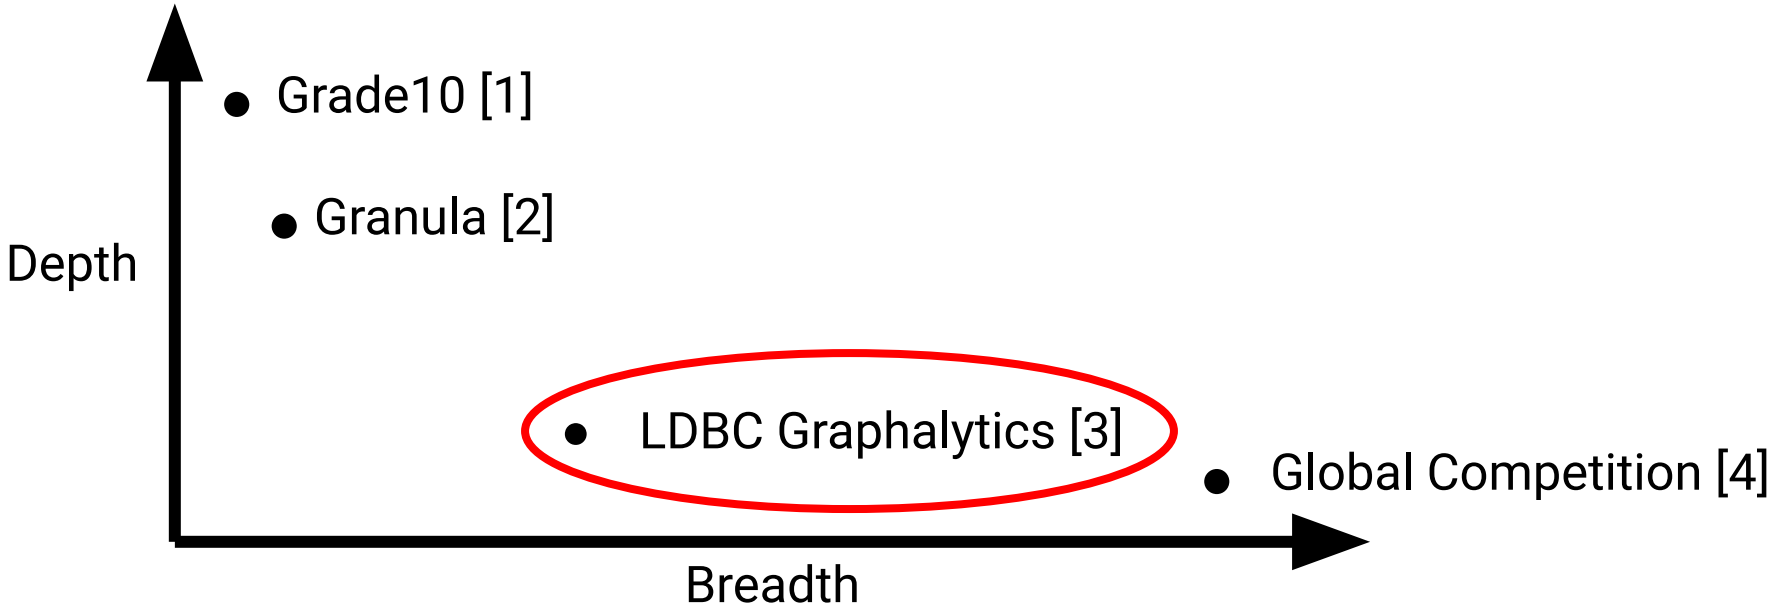

[1] Hegeman, Iosup. Experimental Performance Analysis of Graph Analytics Frameworks. MA thesis. TU Delft, 2018.
[2] Ngai et al., Toward Fine-grained Performance Analysis of Large-scale Graph Processing Platforms. GRADES@SIGMOD/PODS 2017: 8:1-6
[3] Iosup, et al.,. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. PVLDB 9(13): 1317-1328, 2016)
[4] Specification of different Graphalytics Competitions https://graphalytics.org/assets/spec-graphalytics-competitions.pdf, 2018

Comparing Graph Processing Platforms

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, …

Comparing Graph Processing Platforms

How to Compare the Performance of Graph Processing Platforms?

LDBC Graphalytics:

A comprehensive benchmark suite for graph processing across many platforms.

LDBC Graphalytics

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.

LDBC Graphalytics

Main finding:

Performance of graph processing is a non-trivial function of the **PAD** Triangle:

(**Platform**, **Algorithm**, **Dataset**)

+

Hardware, if configurable separately from the Platform

LDBC Graphalytics

Software available at: https://graphalytics.org

- Benchmark core
- Platform drivers (11)
- Data sets (36); S, L, M, XL, 2XL
- Installer
- Documentation

✓ Full support D Directed only U Undirected only ? Incomplete ✗ No support

Platform	BFS	CDLP	LCC	PR	SSSP	WCC
Giraph	✓	✓	✓	✓	✓	✓
GraphX	✓	✓	✓	✓	✓	✓
PowerGraph	✓	✓	✓	✓	✓	✓
OpenG	✓	✓	✓	✓	✓	✓
GraphMat	✓	✓	✓	✓	✓	✓
nvGRAPH	✓	✗	✗	?	✓	✗
Gelly	✓	✓	✓	✓	✓	✓
GraphBLAS	✓	?	✓	D	✓	✓
GraphLab	✓	✓	✓	✗	✗	✗
Neo4j	✓	?	U	D	✓	✓
Gunrock	✓	✗	✗	?	?	✗

The Graphalytics Ecosystem

Graphalytics ecosystem: set of complementary components for understanding graph processing performance.

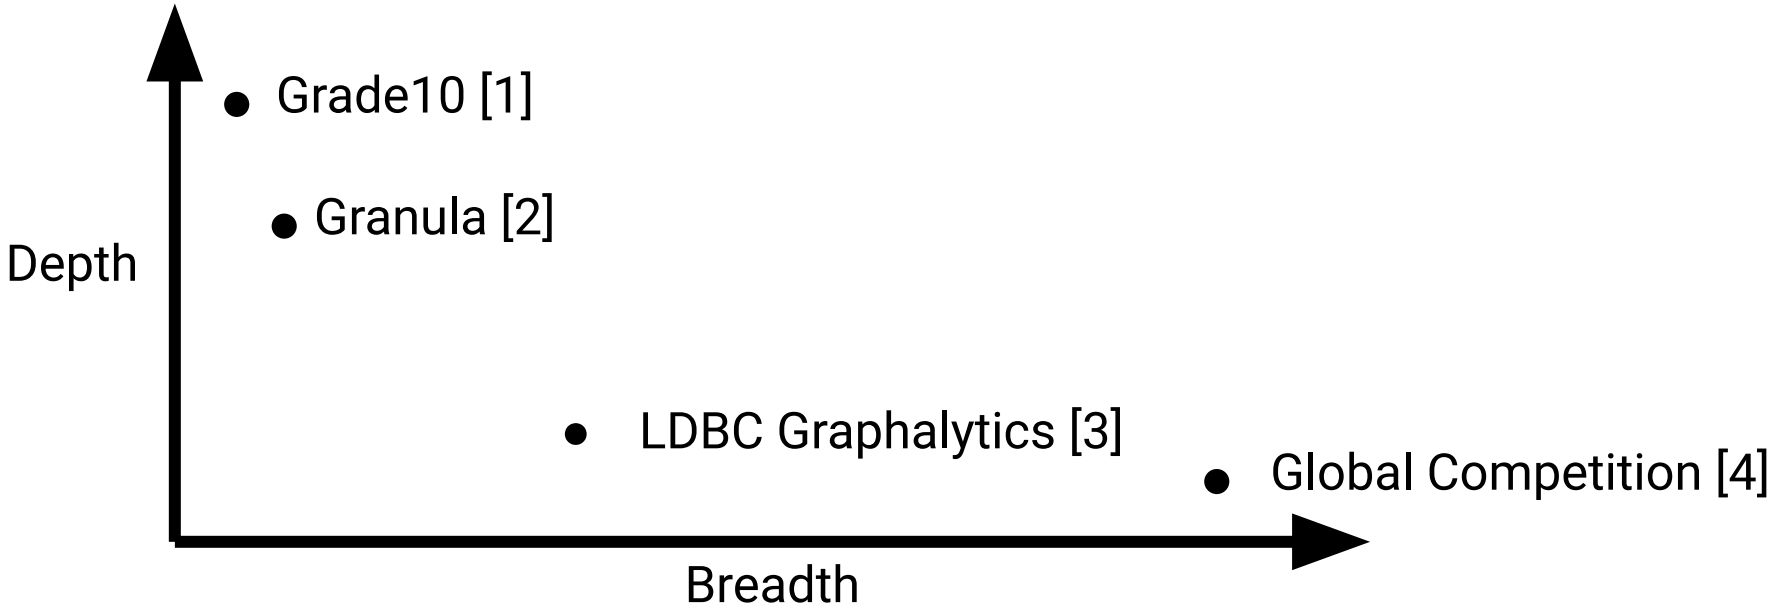

[1] Hegeman, Iosup. Experimental Performance Analysis of Graph Analytics Frameworks. MA thesis. TU Delft, 2018.
[2] Ngai et al., Toward Fine-grained Performance Analysis of Large-scale Graph Processing Platforms. GRADES@SIGMOD/PODS 2017: 8:1-6
[3] Iosup, et al.,. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. PVLDB 9(13): 1317-1328, 2016)
[4] Specification of different Graphalytics Competitions https://graphalytics.org/assets/spec-graphalytics-competitions.pdf, 2018

The Graphalytics Ecosystem

Graphalytics ecosystem: set of complementary components for understanding graph processing performance.

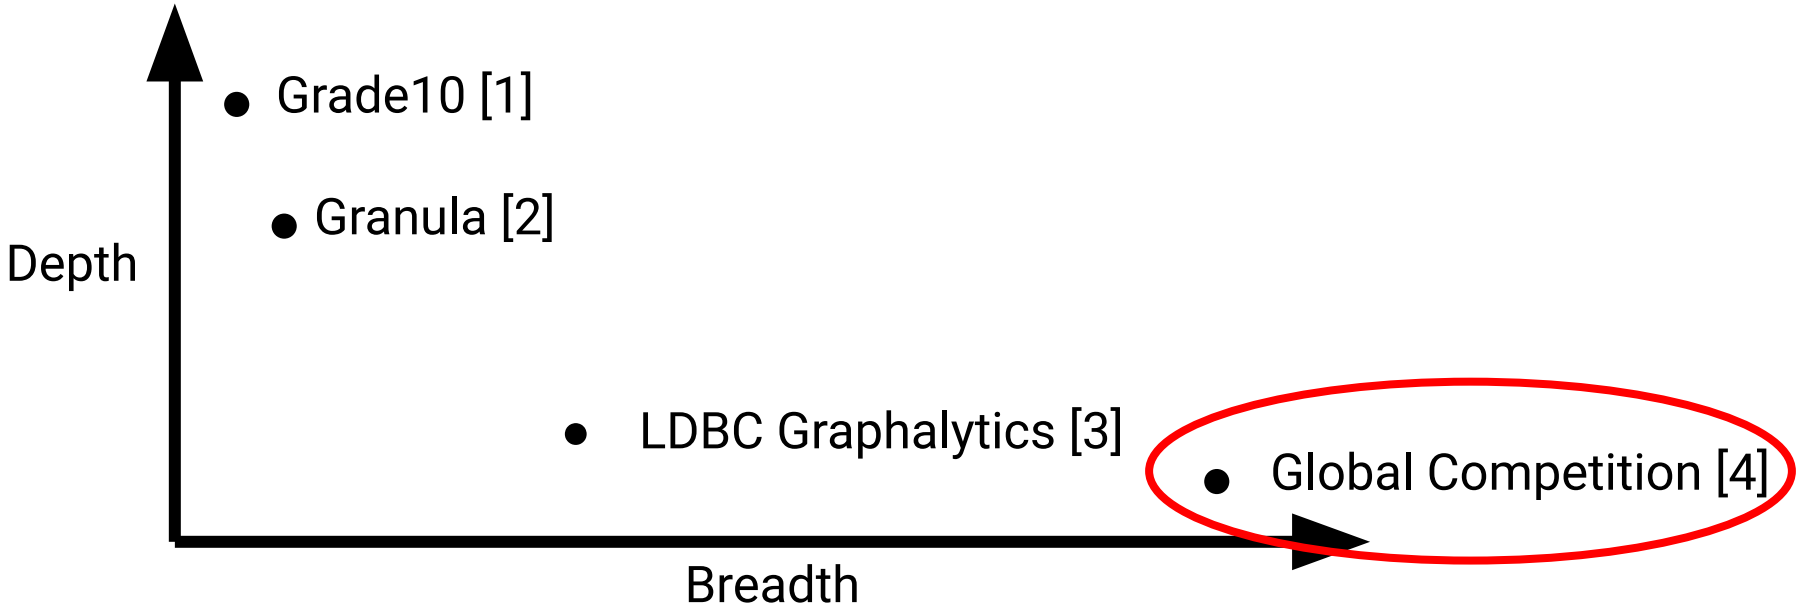

[1] Hegeman, Iosup. Experimental Performance Analysis of Graph Analytics Frameworks. MA thesis. TU Delft, 2018.
[2] Ngai et al., Toward Fine-grained Performance Analysis of Large-scale Graph Processing Platforms. GRADES@SIGMOD/PODS 2017: 8:1-6
[3] Iosup, et al.,. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. PVLDB 9(13): 1317-1328, 2016)
[4] Specification of different Graphalytics Competitions https://graphalytics.org/assets/spec-graphalytics-competitions.pdf, 2018

Graphalytics Global Competition

- Online archive for sharing results
  - https://graphalytics.org/competition
  - Submissions are reviewed by the Graphalytics team
- Systematic and periodic comparison
- Different evaluation metrics
  - Edges and Vertices Per Second (EVPS)
  - Loading Time (TL)
  - Processing Time (PT)
  - Makespan (TM)
- Examples of recent competitions:
  - Google Cloud vs DAS-5
  - CPU vs GPU platforms
  - Various GPUs; GPU platform only competition
- Different scoring methods

Graphalytics Global Competition

Example 'CPU vs GPU platforms' competition
Algorithm: BFS
Metric: Processing Time

Graphalytics Global Competition

Example 'CPU vs GPU platforms' competition
Algorithm: BFS
Metric: Processing Time

Graphalytics Global Competition

Example 'CPU vs GPU platforms' competition
Algorithm: BFS
Metric: Processing Time

Graphalytics Global Competition

Example 'CPU vs GPU platforms' competition
Algorithm: BFS
Metric: Processing Time

Graphalytics Global Competition

Example 'CPU vs GPU platforms' competition
Algorithm: BFS
Metric: Processing Time

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TM)	Datagen-7_5-Fb	Datagen-7_6-Fb	Datagen-7_7-Zf	Dota-League	Graph500-22	Datagen-7_9-Fb
nvgraph @ TitanX	14	0.26 (s) +2	0.30 (s) +2	3.64 (s) +2	0.34 (s) +2	0.55 (s) +2	1.37 (s) +2
nvgraph @ K20	6	1.29 (s) +1	1.31 (s) +1	6.45 (s) +1	1.10 (s) +1	1.34 (s) +1	2.50 (s) +1
nvgraph @ K40	1	1.72 (s) +0	1.94 (s) +0	7.18 (s) +0	1.80 (s) +0	1.75 (s) +0	2.62 (s) +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TM)	Datagen-7_5-Fb	Datagen-7_6-Fb	Datagen-7_7-Zf	Dota-League	Graph500-22	Datagen-7_9-Fb
nvgraph @ TitanX	14	0.26 (s) +2	0.30 (s) +2	3.64 (s) +2	0.34 (s) +2	0.55 (s) +2	1.37 (s) +2
nvgraph @ K20	6	1.29 (s) +1	1.31 (s) +1	6.45 (s) +1	1.10 (s) +1	1.34 (s) +1	2.50 (s) +1
nvgraph @ K40	1	1.72 (s) +0	1.94 (s) +0	7.18 (s) +0	1.80 (s) +0	1.75 (s) +0	2.62 (s) +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TM)	Datagen-7_5-Fb	Datagen-7_6-Fb	Datagen-7_7-Zf	Dota-League	Graph500-22	Datagen-7_9-Fb
nvgraph @ TitanX	14	0.26 (s) +2	0.30 (s) +2	3.64 (s) +2	0.34 (s) +2	0.55 (s) +2	1.37 (s) +2
nvgraph @ K20	6	1.29 (s) +1	1.31 (s) +1	6.45 (s) +1	1.10 (s) +1	1.34 (s) +1	2.50 (s) +1
nvgraph @ K40	1	1.72 (s) +0	1.94 (s) +0	7.18 (s) +0	1.80 (s) +0	1.75 (s) +0	2.62 (s) +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TM)	Datagen-7_5-Fb	Datagen-7_6-Fb	Datagen-7_7-Zf	Dota-League	Graph500-22	Datagen-7_9-Fb
nvgraph @ TitanX	14	0.26 (s) +2 +2	0.30 (s) +2 +2	3.64 (s) +2 +2	0.34 (s) +2 +2	0.55 (s) +2 +2	1.37 (s) +2 +2
nvgraph @ K20	6	1.29 (s) +1 +1	1.31 (s) +1 +1	6.45 (s) +1 +1	1.10 (s) +1 +1	1.34 (s) +1 +1	2.50 (s) +1 +1
nvgraph @ K40	1	1.72 (s) +0 +0	1.94 (s) +0 +0	7.18 (s) +0 +0	1.80 (s) +0 +0	1.75 (s) +0 +0	2.62 (s) +0 +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TM)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
nvgraph @ TitanX	8	8.11 (s) +2 +2	12.24 (s) +1 +1	237.48 (s) +2 +2	9.29 (s) +1 +1	250.67 (s) +2 +2	- +0 +0
nvgraph @ K40	7	9.47 (s) +1 +1	11.45 (s) +2 +2	305.96 (s) +1 +1	7.78 (s) +2 +2	307.90 (s) +1 +1	- +0 +0
nvgraph @ K20	0	- +0 +0	- +0 +0	- +0 +0	- +0 +0	- +0 +0	- +0 +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TM)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
nvgraph @ TitanX	8	8.11 (s) +2 +2	12.24 (s) +1 +1	237.48 (s) +2 +2	9.29 (s) +1 +1	250.67 (s) +2 +2	- +0 +0
nvgraph @ K40	7	9.47 (s) +1 +1	11.45 (s) +2 +2	305.96 (s) +1 +1	7.78 (s) +2 +2	307.90 (s) +1 +1	- +0 +0
nvgraph @ K20	0	- +0 +0	- +0 +0	- +0 +0	- +0 +0	- +0 +0	- +0 +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TP)	Datagen-7_5-Fb	Datagen-7_6-Fb	Datagen-7_7-Zf	Dota-League	Datagen-7_8-Zf	Graph500-22	Datagen-7_9-Fb
nvgraph @ K40	10	0.31 (s) +1 +1	0.40 (s) +1 +1	1.28 (s) +2 +2	0.18 (s) +1 +1	2.00 (s) +2 +2	0.04 (s) +2 +2	0.95 (s) +1 +1
nvgraph @ TitanX	8	0.10 (s) +2 +2	0.12 (s) +2 +2	3.37 (s) +0 +0	0.09 (s) +2 +2	5.68 (s) +0 +0	0.21 (s) +0 +0	0.46 (s) +2 +2
nvgraph @ K20	3	0.38 (s) +0 +0	0.48 (s) +0 +0	1.50 (s) +1 +1	0.23 (s) +0 +0	2.33 (s) +1 +1	0.05 (s) +1 +1	1.10 (s) +0 +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TP)	Datagen-7_5-Fb	Datagen-7_6-Fb	Datagen-7_7-Zf	Dota-League	Datagen-7_8-Zf	Graph500-22	Datagen-7_9-Fb
nvgraph @ K40	10	0.31 (s) +1 +1	0.40 (s) +1 +1	1.28 (s) +2 +2	0.18 (s) +1 +1	2.00 (s) +2 +2	0.04 (s) +2 +2	0.95 (s) +1 +1
nvgraph @ TitanX	8	0.10 (s) +2 +2	0.12 (s) +2 +2	3.37 (s) +0 +0	0.09 (s) +2 +2	5.68 (s) +0 +0	0.21 (s) +0 +0	0.46 (s) +2 +2
nvgraph @ K20	3	0.38 (s) +0 +0	0.48 (s) +0 +0	1.50 (s) +1 +1	0.23 (s) +0 +0	2.33 (s) +1 +1	0.05 (s) +1 +1	1.10 (s) +0 +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TP)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
nvgraph @ TitanX	10	2.93 (s) +2 +2	3.89 (s) +2 +2	209.50 (s) +2 +2	0.40 (s) +2 +2	162.36 (s) +2 +2	- +0 +0
nvgraph @ K40	5	4.02 (s) +1 +1	5.13 (s) +1 +1	278.62 (s) +1 +1	0.54 (s) +1 +1	218.09 (s) +1 +1	- +0 +0
nvgraph @ K20	0	- +0 +0	- +0 +0	- +0 +0	- +0 +0	- +0 +0	- +0 +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TP)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
nvgraph @ TitanX	10	2.93 (s) +2 +2	3.89 (s) +2 +2	209.50 (s) +2 +2	0.40 (s) +2 +2	162.36 (s) +2 +2	- +0 +0
nvgraph @ K40	5	4.02 (s) +1 +1	5.13 (s) +1 +1	278.62 (s) +1 +1	0.54 (s) +1 +1	218.09 (s) +1 +1	- +0 +0
nvgraph @ K20	0	- +0 +0	- +0 +0	- +0 +0	- +0 +0	- +0 +0	- +0 +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TP)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
nvgraph @ TitanX	10	2.93 (s) +2 +2	3.89 (s) +2 +2	209.50 (s) +2 +2	0.40 (s) +2 +2	162.36 (s) +2 +2	- +0 +0
nvgraph @ K40	5	4.02 (s) +1 +1	5.13 (s) +1 +1	278.62 (s) +1 +1	0.54 (s) +1 +1	218.09 (s) +1 +1	- +0 +0
nvgraph @ K20	0	- +0 +0	- +0 +0	- +0 +0	- +0 +0	- +0 +0	- +0 +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TP)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
nvgraph @ TitanX	12	0.07 (s) +2 +2	0.09 (s) +2 +2	0.45 (s) +2 +2	0.16 (s) +2 +2	0.60 (s) +2 +2	0.20 (s) +2 +2
nvgraph @ K40	6	0.22 (s) +1 +1	0.31 (s) +1 +1	0.72 (s) +1 +1	0.41 (s) +1 +1	0.92 (s) +1 +1	0.56 (s) +1 +1
nvgraph @ K20	0	0.28 (s) +0 +0	0.37 (s) +0 +0	- +0 +0	0.51 (s) +0 +0	- +0 +0	- +0 +0

Graphalytics Global Competition

Example 'Various GPUs' competition
Algorithm: PR - BFS
Metric: Makespan - Processing Time
Dataset scale: Small - Large

System name	Total score (TP)	Datagen-8_5-Fb	Datagen-8_6-Fb	Datagen-8_7-Zf	Graph500-25	Datagen-8_8-Zf	Datagen-8_9-Fb
nvgraph @ TitanX	12	0.07 (s) +2 +2	0.09 (s) +2 +2	0.45 (s) +2 +2	0.16 (s) +2 +2	0.60 (s) +2 +2	0.20 (s) +2 +2
nvgraph @ K40	6	0.22 (s) +1 +1	0.31 (s) +1 +1	0.72 (s) +1 +1	0.41 (s) +1 +1	0.92 (s) +1 +1	0.56 (s) +1 +1
nvgraph @ K20	0	0.28 (s) +0 +0	0.37 (s) +0 +0	- +0 +0	0.51 (s) +0 +0	- +0 +0	- +0 +0

Graphalytics Global Competition

Full results & competition reports available:
https://graphalytics.org/competition

Graphalytics competition report: 2018 GPU vs CPU platforms

OCTOBER 2018 - GRAPHALYTICS.ORG

This report describes the "*2018 GPU vs CPU platforms*" (tournament and relative-performance) competitions listed on the Graphalytics competitions page. Questions or feedback can be e-mailed to the Graphalytics team [1].

1 OVERVIEW

The following table lists the two competitions that are described in this report:

Table 1. List of competitions described this report.

Name	Description	Type
2018 GPU vs CPU platforms	S/M/L/XL	Tournament
2018 GPU vs CPU platforms	S/M/L/XL	Relative performance

The nvGRAPH library is part of the CUDA toolkit developed by NVIDIA and is the first GPU library in the Graphalytics benchmark. Table 2 lists the platforms that are being used in this competition. By default, nvGRAPH implements BFS, SSSP and PR, and provides example code of these algorithms in the nvGRAPH documentation. Furthermore, the results of nvGRAPH's PR were removed as it failed to pass the Graphalytics validation test.

Gunrock implements BFS, PR and SSSP. However, only BFS is benchmarked in this competition, as it produced correct (validated) results. PR lacks parameters to set the number of iterations and damping, while SSSP does not support floating point units.

The CPU platforms support all the Graphalytics standard algorithms.

Rationale of single/multi node CPU platforms: Note that the GPU platforms were benchmarked on a single node. For CPU platforms, the results for single- and multi-node benchmark results are added for GraphMat to give a better understanding of how GPU platforms rank. In addition, the benchmark results for PowerGraph (single node) are added as well.

Table 2. List of platforms used in the "2018 GPU vs CPU platforms" competitions

Platform	Core-Version	Driver-Version	Platform-Version	#Algorithms
nvGRAPH	1.0.0	0.2-SNAPSHOT	Toolkit v10.0.130	2/6
Gunrock	1.0.0	0.2-SNAPSHOT	Commit ea18455	1/6
Giraph	0.9.0/1.0.0	0.2-SNAPSHOT	1.2.0-hadoop2	6/6
GraphMat	0.9.0/1.0.0	0.2-SNAPSHOT	May 2017 (a892c36)	6/6
PowerGraph @ single node	1.0.0	0.2-SNAPSHOT	2.2	6/6
GraphMat @ single node	1.0.0	0.2-SNAPSHOT	May 2017 (a892c36)	6/6

[1] https://graphalytics.org/about

1

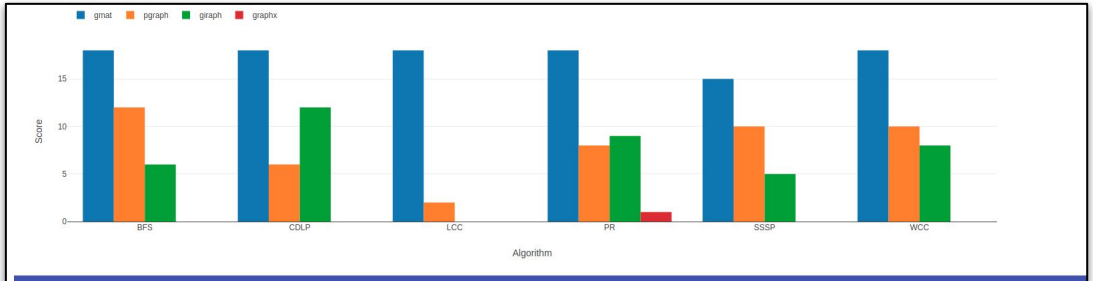

Rank	System name	Total score	BFS	CDLP	LCC	PR	SSSP	WCC
No. 1	gmat	105	18	18	18	18	15	18
No. 2	pgraph	48	12	6	2	8	10	10
No. 3	giraph	40	6	12	0	9	5	8
No. 4	graphx	1	0	0	0	1	0	0

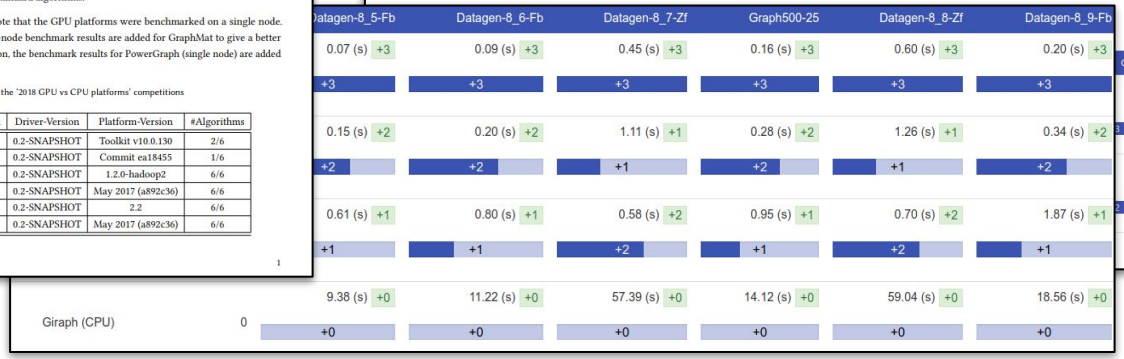

Graphalytics 2.0

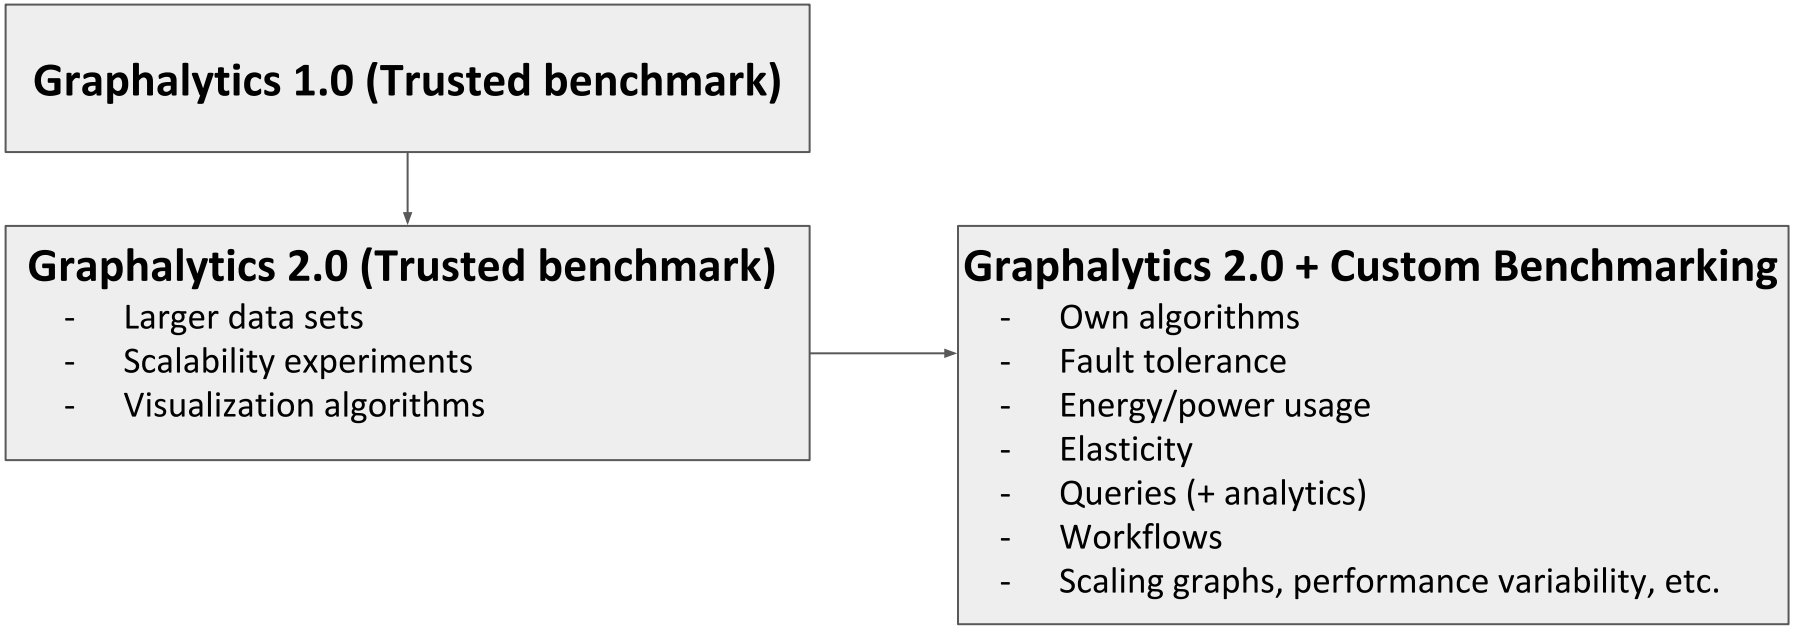

Take home message

- The **Graphalytics ecosystem** provides **breadth** and **depth** in understanding graph processing performance.
- View & submit benchmark results @ **Graphalytics Global Competition**

graphalytics.org

[1] Hegeman, Iosup. Experimental Performance Analysis of Graph Analytics Frameworks. MA thesis. TU Delft, 2018.
[2] Ngai et al., Toward Fine-grained Performance Analysis of Large-scale Graph Processing Platforms. GRADES@SIGMOD/PODS 2017: 8:1-6
[3] Iosup, et al.,. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. PVLDB 9(13): 1317-1328, 2016)
[4] Specification of different Graphalytics Competitions https://graphalytics.org/assets/spec-graphalytics-competitions.pdf, 2018