

Generous donation from:



Co-sponsored by:



COMMIT/



Graphalytics:

A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms



Tim Hegeman, Wing-Lung Ngai, and Stijn Heldens.



Presentation developed jointly with Ana Lucia Varbanescu.



Several slides developed jointly with Yong Guo.



Prof. dr. ir. Alexandru Iosup

Massivizing Computer Systems

Co-authored by LDBC team:

Arnau Prat-Pérez, Thomas Manhardt, Siegfried Depner, Hassan Chafi, Mihai Capotă, Narayanan Sundaram, Michael Anderson, Ilie Gabriel Tănase, Yinglong Xia, Lifeng Nai, Peter Boncz



VU Amsterdam / TU Delft – the Netherlands – Europe



founded 10th century
pop: 850,000



founded 1880
pop: 23,500

Amsterdam



Delft

founded 1842
pop: 19,500



founded 13th century
pop: 100,000



pop: 16.5 M



Walldorf, Germany

GraphsComp in Academic Publications

Title Keywords in Computer Systems Conferences (CCGRID, CLOUD, Cluster, HPDC, ICPP, IPDPS, NSDI, OSDI, SC, SIGMETRICS, SoCC, SOSp,) and Journals (CCPE, FGCS, JPDC, TPDS)

Rank	2016	Change	2015	Change	2014	Change	2013	Change	2012	Change
1	cloud	-	cloud	0	cloud	0	cloud	0	cloud	0
2	data	-	data	0	data	0	data	0	data	0
29	graph	-	dynamic	-3	management	-1	architecture	+1	framework	-2
30	machine	-	machine	0	architecture	+1	analysis	-3	mapreduce	+4
31	virtual	-	architecture	+1	mapreduce	+2	center	+4	core	+6
32	architecture	-	graph	-3	machine	-2	machine	0	model	-10
33	time	-	mapreduce	+7	time	+1	graph	+1	center	-2
34	approach	-	time	-1	graph	-2	mapreduce	-3	virtual	-10
35	center	-	center	0	center	0	simulation	+1	multicore	+7
36	optimization	-	approach	-2	simulation	+1	heterogeneous	+2	graph	-3

Warning: Linear regressions may be deceiving.

Graphs Are at the Core of Our Society: The LinkedIn Example

The State of **LinkedIn**

A very good resource for matchmaking workforce and prospective employers


Vital for your company's life,
as your Head of HR would tell you

Vital for the prospective employees


Tens of “specialized LinkedIns”: medical, mil, edu, science, ...

~~1~~ **50,000,000**
registered members (Q1 '12)

LinkedIn's Service Analysis



LinkedIn's Service Analysis



How to do Graph Analysis? Graph Processing @large

Linked **in**



A Graph Processing Platform



friendster

XFIRE™

Graph Processing Platforms

Which platforms perform well?

What to tune?

What to re-design?

ORACLE

Neo4j
the graph database

PROJECT PE

ACHIE
GRAPH


Lab

GraphX

Graph Processing Platforms

ORACLE PGX

Intel Graphmat



IBM System G

TOTEM

Benchmark!



GraphDB

Trinity



What Is the Performance of Graph Processing Platforms?

**Metrics
Diversity**

**Graph
Diversity**

**Algorithm
Diversity**

- Graph500
 - Single application (BFS), Single class of synthetic datasets. @ISC16: future diversification.
- Few existing platform-centric comparative studies
 - Prove the superiority of a given system, limited set of metrics
- GreenGraph500, GraphBench, XGDBench
 - Issues with representativeness, systems covered, metrics, ...

What Is the Performance of Graph Processing Platforms?

Metrics
Diversity

Graph
Diversity

Algorithm
Diversity



Graphalytics = comprehensive benchmarking suite for graph processing across many platforms

<http://ldbncouncil.org/ldbnc-graphalytics>

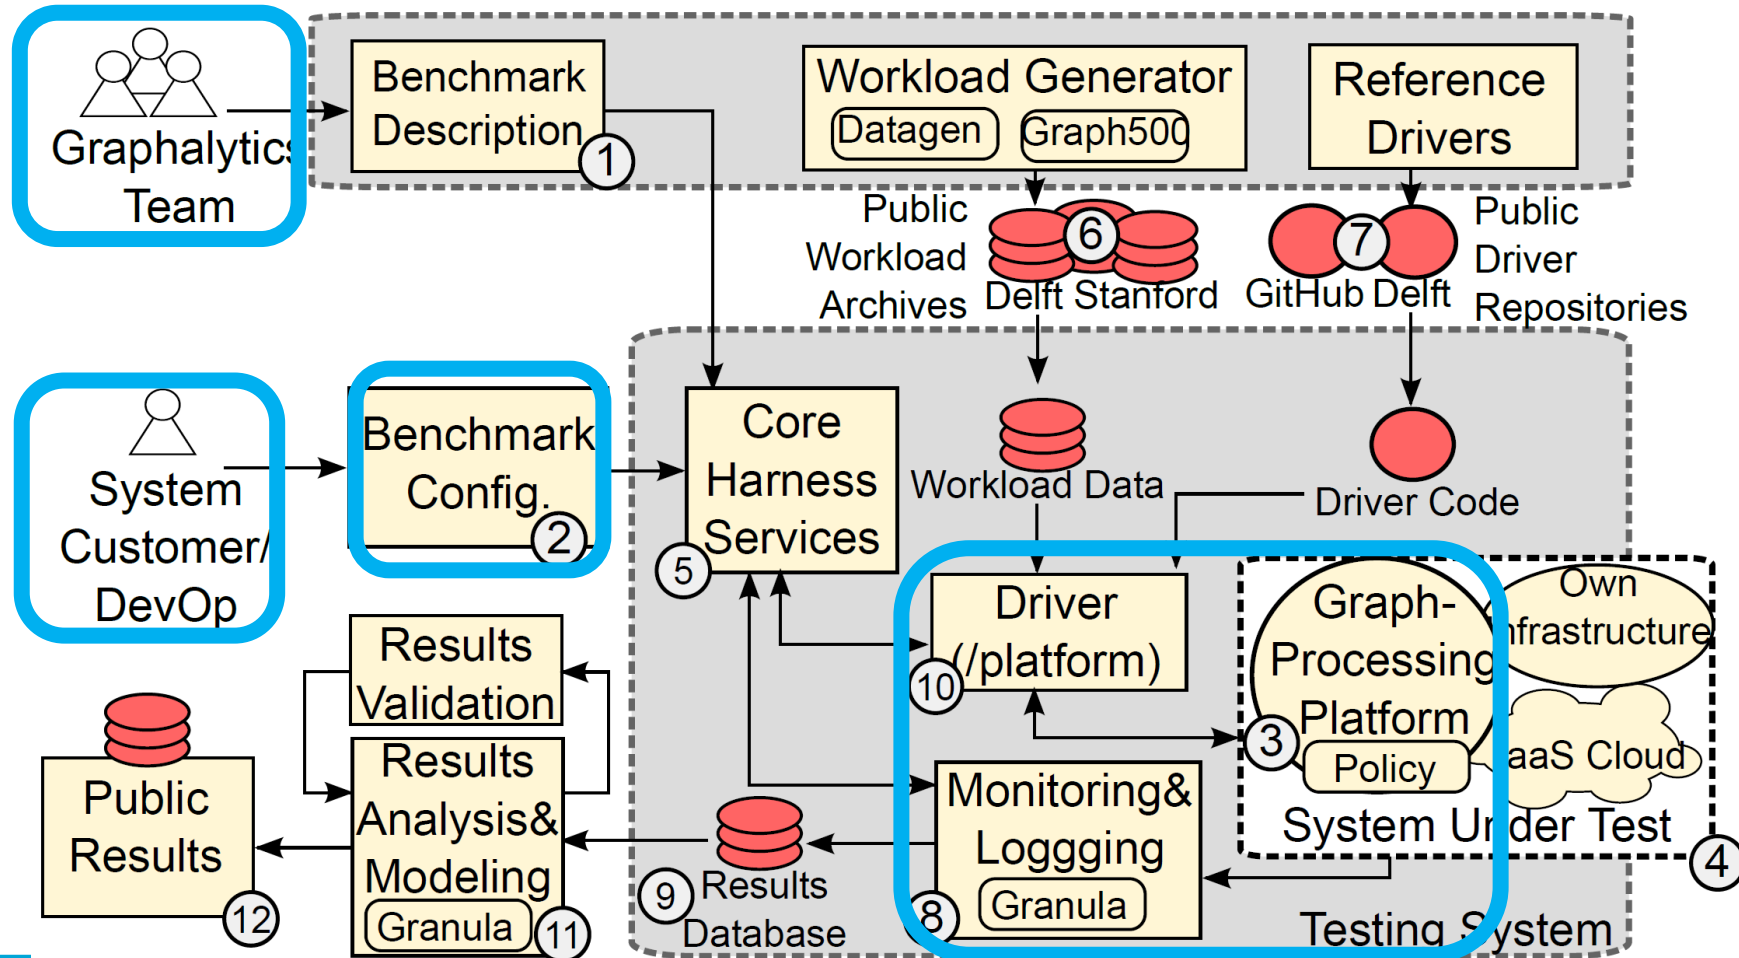
<http://graphalytics.ewi.tudelft.nl/>

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
- Extended toolset for **manual choke-point analysis**
- **Enables comparison of many platforms, community-driven and industrial**




Graphalytics = Benchmarking Harness



Iosup et al. LDBC Graphalytics: A Benchmark for Large Scale Graph Analysis on Parallel and Distributed Platform, PVLDB'16.

Graphalytics = Representative Classes of Algorithms and Datasets



- 2-stage selection process of algorithms and datasets

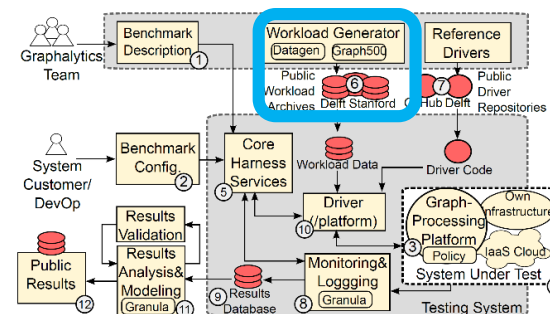
Class	Examples	%
Graph Statistics	Diameter, Local Clust. Coeff., PageRank	20
Graph Traversal	BFS, SSSP, DFS	50
Connected Comp.	Reachability, BiCC, Weakly CC	10
Community Detection	Clustering, Nearest Neighbor, Community Detection w Label Propagation	5
Other	Sampling, Partitioning	<15

+ property/weighted graphs: Single-Source Shortest Paths (~35%)

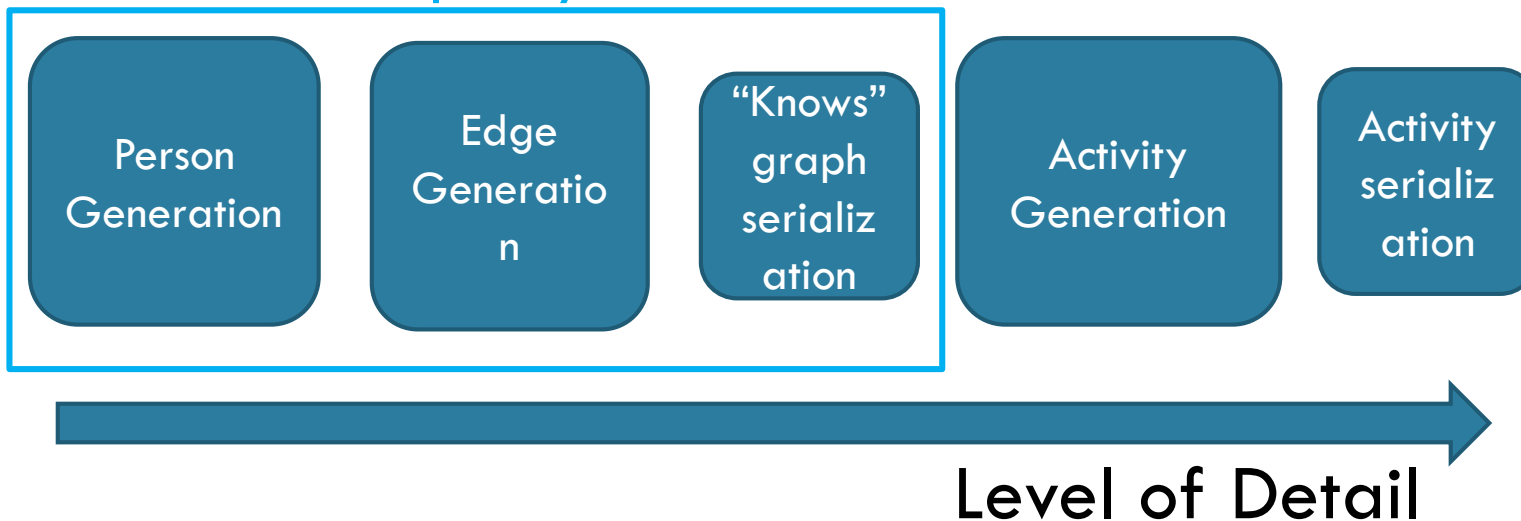


Graphalytics = Distributed Graph Generation with DATAGEN

- Rich set of configurations
- More diverse degree distribution than Graph500
- Realistic clustering coefficient and assortativity



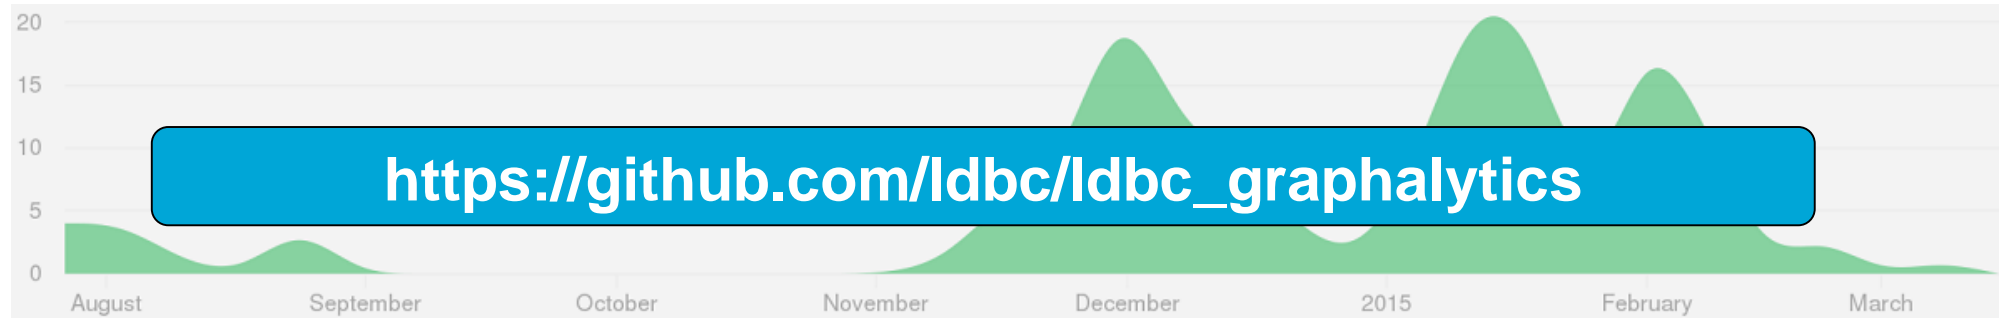
Graphalytics



Graphalytics = Diverse Set of Automated Experiments

Category	Experiment	Algo.	Data	Nodes/ Threads	Metrics
Baseline	Dataset variety	BFS,PR	All	1	Run, norm.
	Algorithm variety	All	R4(S), D300(L)	1	Runtime
Scalability	Vertical vs. horiz.	BFS, PR	D300(L), D1000(XL)	1—16/1—32	Runtime, S
	Weak vs. strong	BFS, PR	G22(S)— G26(XL)	1—16/1—32	Runtime, S
Robustness	Stress test	BFS	All	1	SLA met
	Variability	BFS	D300(L), D1000(L)	1/16	CV
Self-Test	Time to run/part	--	Datagen	1—16	Runtime

Graphalytics = Modern Software Engineering Process



Graphalytics code reviews

Internal release to LDBC partners (first, Feb 2015; last, Feb 2016)

Public release, announced first through LDBC (Apr 2015)

First full benchmark specification, LDBC criteria (Q1 2016)

Jenkins continuous integration server

SonarQube software quality analyzer

Ongoing Activity in the Graphalytics Team (2016-2017)



1. A public, curated database of rated graph-processing platforms

- Demo follows in next presentation



2. Grade10: systematic analysis of performance bottlenecks


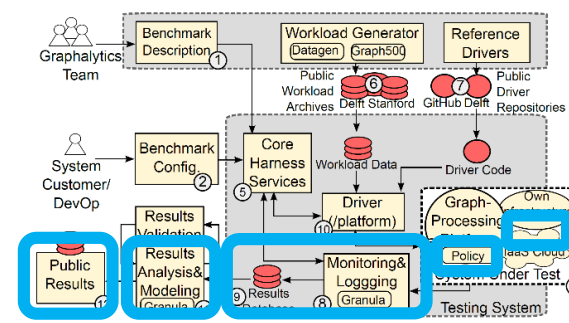


3. Granula: process for modeling, modeling, archiving, and sharing performance results for graph-processing platforms



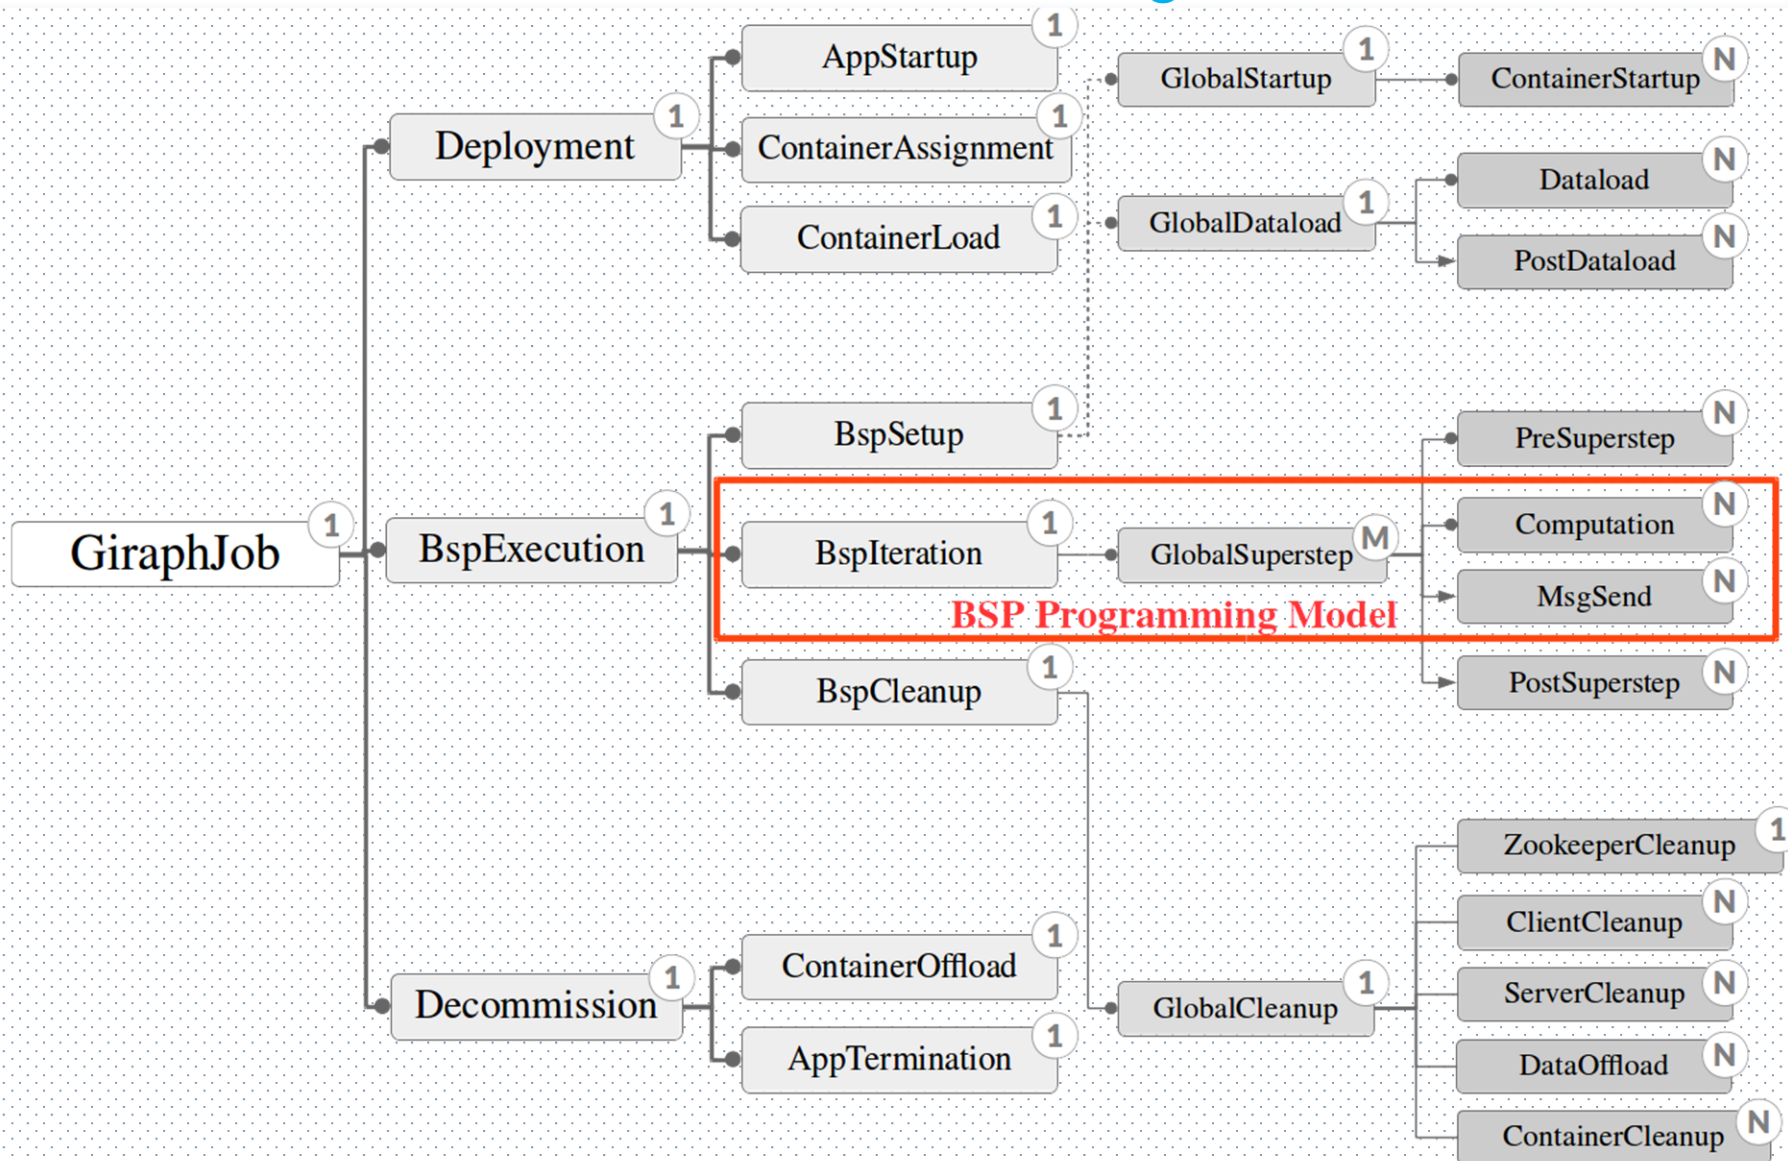
4. Release of full-fledged LDBC Graphalytics benchmark

Graphalytics = Portable Performance Analysis with Granula



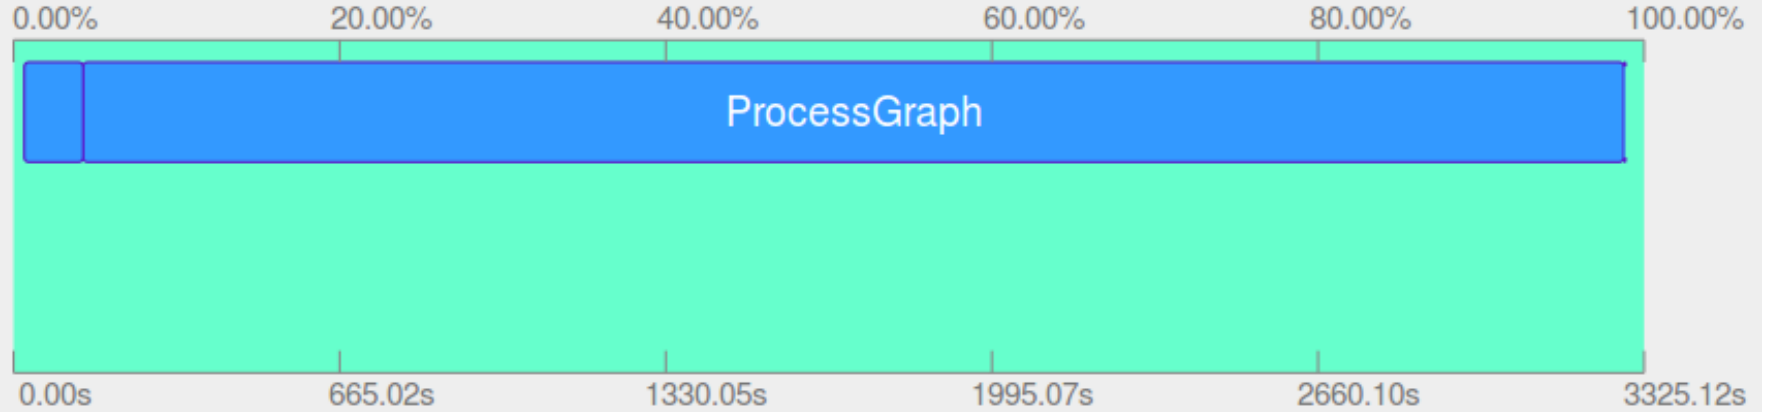
Minimal code invasion + automated data collection at runtime + portable archive (+ web UI) → portable bottleneck analysis

Incremental Performance Modelling with Granula




Performance Monitoring, Archiving, Visualization with Granula

Explore HERE! >




Environment Data




Giraph - CDLP on LDBC-1000, 8 nodes Cpu Time

Performance Visualization, Analysis with Granula


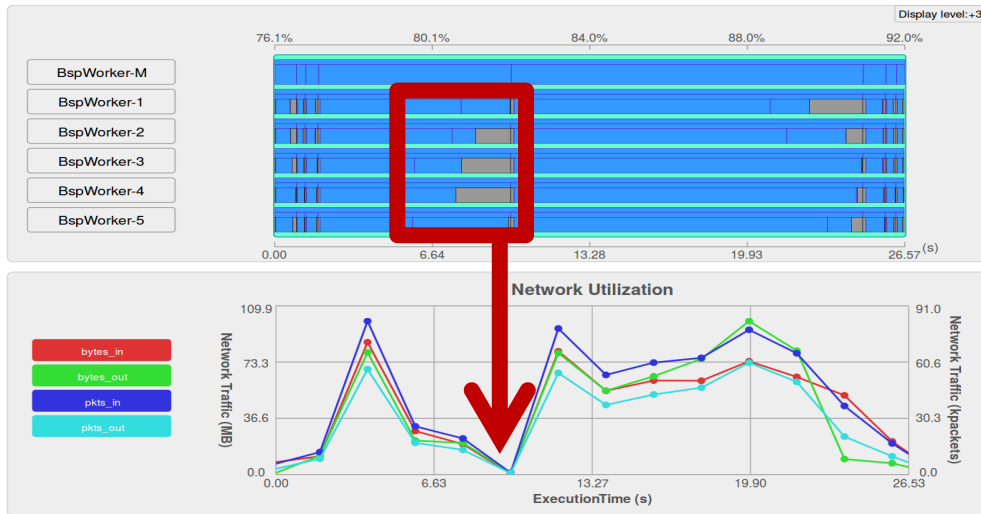


Giraph - BFS on LDBC-1000, 5 nodes



Grade10: Performance Bottleneck Identification

Performance analysis is **time-consuming** and **expertise-driven**.
Grade10 analyses Granula & resource utilization data for you.

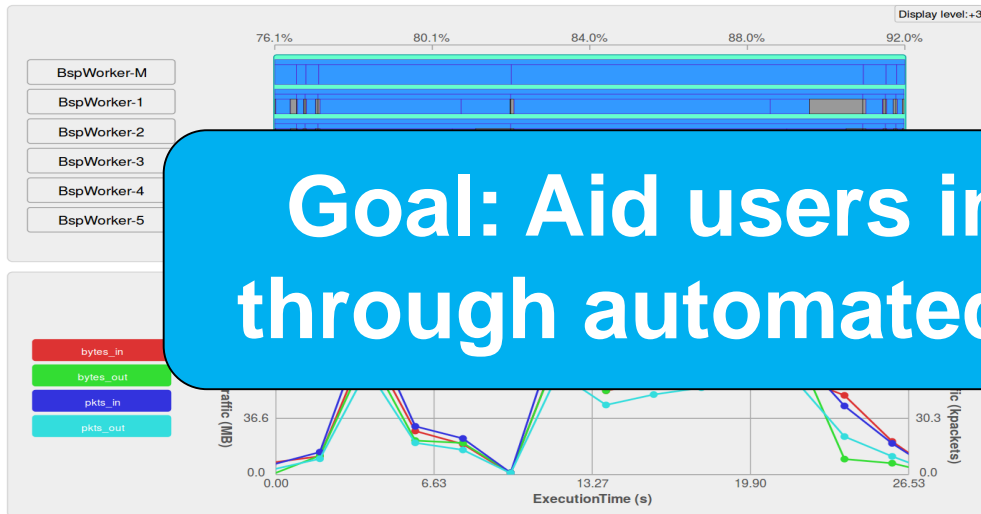


Possible performance bottlenecks:

- 20% slowdown due to imbalance in 'Computation' phase
- HW resource bottlenecks of 'GlobalSuperstep': CPU 60%, network 30%, none 10%

Grade10: Performance Bottleneck Identification

Performance analysis is **time-consuming** and **expertise-driven**.
Grade10 analyses Granula & resource utilization data for you.



Possible performance bottlenecks:

Goal: Aid users in understanding performance through automated analysis of performance data

CPU 60%, network 30%, none 10%

Grade10: Performance Bottleneck Identification

Possible future directions:

1. Support **performance regression tests** by identifying shifts in bottlenecks
2. Identify **platform-wide bottlenecks** through systematic evaluation of Graphalytics results
3. Integrate **low-level performance data**, including HW performance counters, tracing data

Full Benchmark: 4 Types of Benchmarks

1. Test benchmark / fire drill

2. Standard benchmark

- cost-efficiency*, performance

3. Full benchmark

- scalability, robustness

4. Custom benchmark

- specialized analysis, based on Granula and Grade10



A public, curated DB of
rated graph-processing
platforms

Graphalytics Roadmap

Date	Release	Competition	Activities
2017-01-30	v0.2.8	Beta Competition: R2	Refine standard benchmark definition + cost-efficiency + performance
2017-03-13	v0.2.9	Beta Competition: R3	Refine system specification, cost model
2017-04-10	v0.2.10	Beta Competition: R3	Refine full benchmark definition + scalability + robustness
2017-05-08	v0.2.11	Beta Competition: R3	Refine competition, auditing Rules
2017-06-05	v0.2.12	Beta Competition: R3	[reserved slot]
2017-06-19	v1.0.0	2017, Edition 1: Completed	Internal participation
2017-06-26	v1.0.0	2017, Edition 2: Started	Global participation

Graphalytics, in the future

An LDBC benchmark*
Advanced benchmarking harness
Diverse real and synthetic datasets
Many classes of algorithms
Granula, Grade10 for bottleneck analysis
Modern software engineering practices
Supports many platforms
Enables comparison of
community-driven and industrial systems
Public, curated DB of rated systems

+ more data generation
+ deeper performance metrics
+ bottleneck analysis

LDBC  *The graph & RDF
benchmark reference*




UNIVERSITEIT
VAN AMSTERDAM



  HUAWEI

ORACLE[®]