

Generous donation from:



Co-sponsored by:



COMMIT/

Graphalytics team hosted by:




Graphalytics:

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



Tech Leads: Wing-Lung Ngai and Tim Hegeman.



Collaborators: Ana Lucia Varbanescu and Stijn Heldens.



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



Munich,
Germany

GraphsComp in Academic Publications

Title Keywords in Computer Systems Conferences (CCGRID, CLOUD, Cluster, HPDC, ICPP, IPDPS, NSDI, OSDI, SC, SIGMETRICS, SoCC, SOSIP) 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



Graph Processing Platforms

Which platforms perform well?

**What to tune?
What to re-design?**

ORACLE

Neo4j
the graph database

PROJECT PE

A C H E
R A P H


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.org/>

Graphalytics, in a nutshell

- An LDBC benchmark <http://ldbcouncil.org/ldbc-graphalytics>
- 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

<http://ldbcouncil.org/ldbc-graphalytics>



[Iosup et al., VLDB'16] [Guo et al., CCGRID'15]
[Guo et al., IPDPS'14]

<http://graphalytics.org>

Graphalytics Software: 4 Types of Benchmarks

1. **Test benchmark / fire drill**
2. **Standard benchmark**
 - cost-efficiency*, performance
3. **Full benchmark**
 - scalability, robustness
4. **Custom benchmark**
 - analysis, based on our tools Granula and Grade10



Benchmark +
Global Competition
=
A public, curated DB of
rated graph-processing
platforms

Ongoing Activity in the Graphalytics Team (2016-2017)



1. Benchmark + Global Competition

[Iosup et al., VLDB'16]++

- Details follow in next presentation by Tech Lead Wing Lung Ngai



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

- Presented in Walldorf, February 2017


[Ngai et al., GRADES'17]




3. Grade10: automated bottleneck detection and performance-issue identification

4. (end-2017) Release of Graphalytics Global Competition (2 new hires)

Granula: Performance Modeling, Visualization, Analysis





Giraph - BFS on LDBC-1000, 5 nodes



Grade10: Performance Bottleneck Identification

Analytical modeling is **time-consuming**. Profiling (aggregating) and full tracing are **data-intensive**. All are **expertise-driven**.

Grade10 analyses Granula and resource utilization data for you.

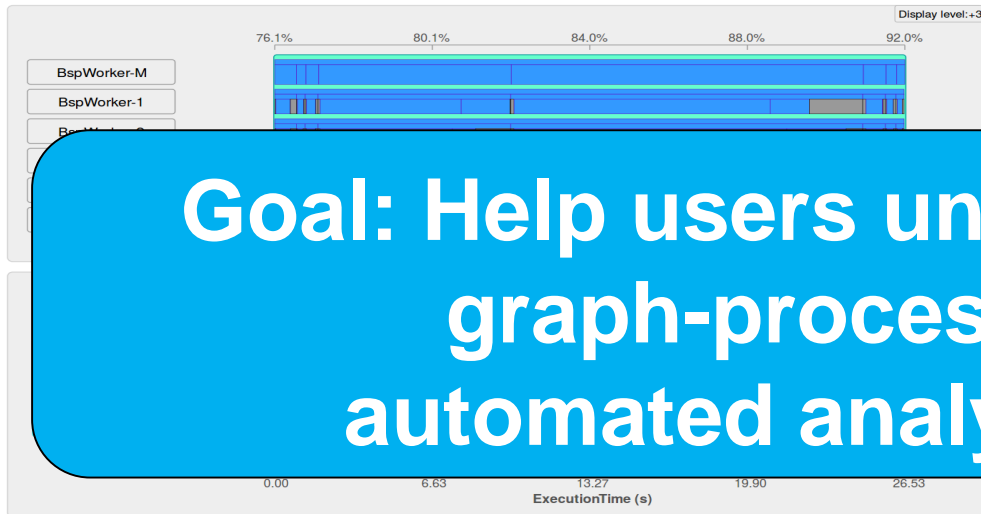


Possible performance bottlenecks:

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

Grade10: Performance Bottleneck Identification

Analytical modeling is **time-consuming**. Profiling (aggregating) and full tracing are **data-intensive**. All are **expertise-driven**.
Grade10 analyses Granula and resource utilization data for you.




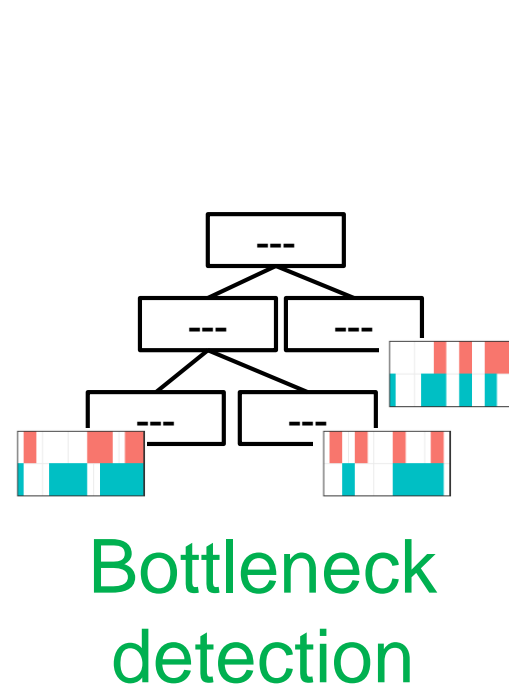

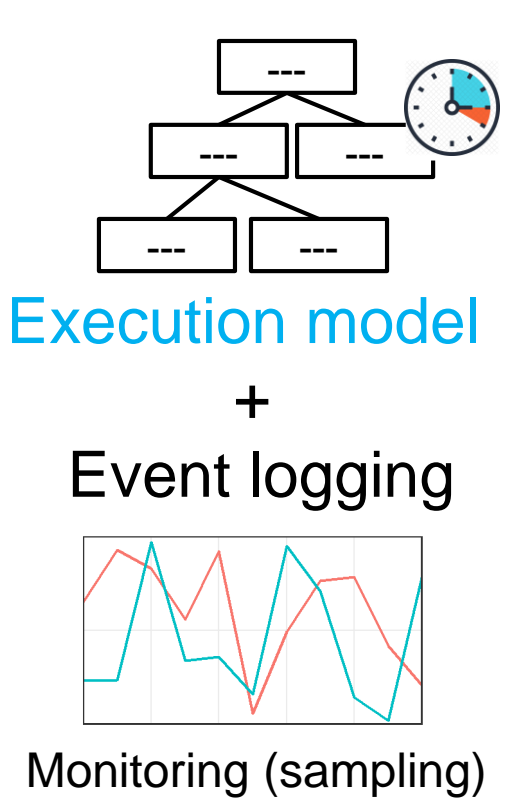
Possible performance bottlenecks:

Goal: Help users understand the performance of graph-processing systems through automated analysis of performance data


Grade10: Automated Bottleneck Detection and Performance Issue Identification



System under test



Preliminary Result: Analysing a Giraph Job





WorkerSuperstep



**CPU usage < 32 cores
(100%), so no bottleneck**

... yet

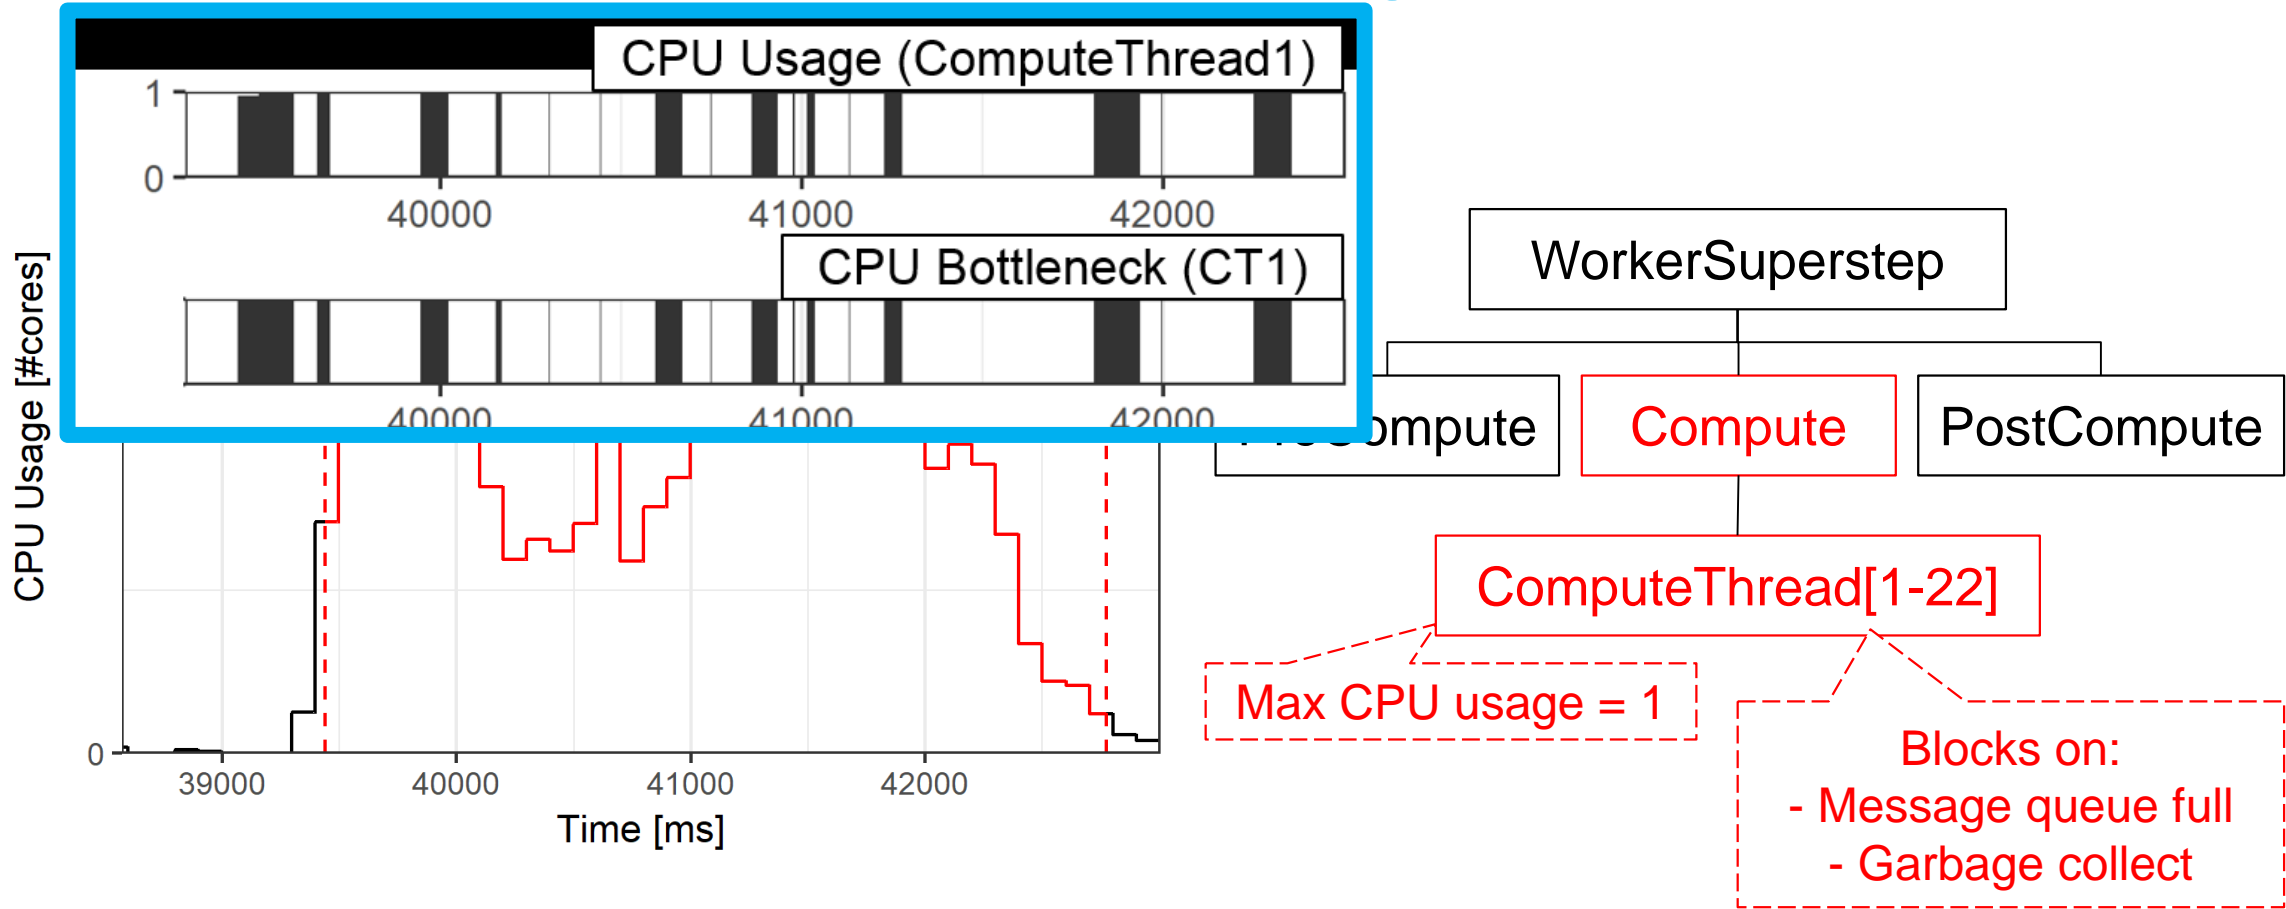
Preliminary Result: Analysing a Giraph Job



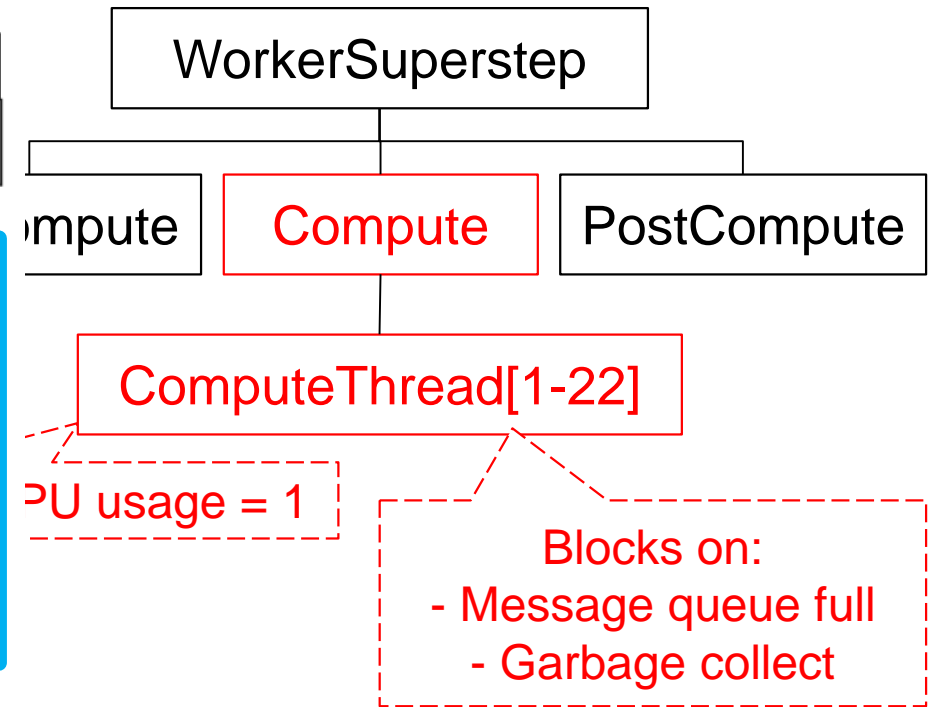
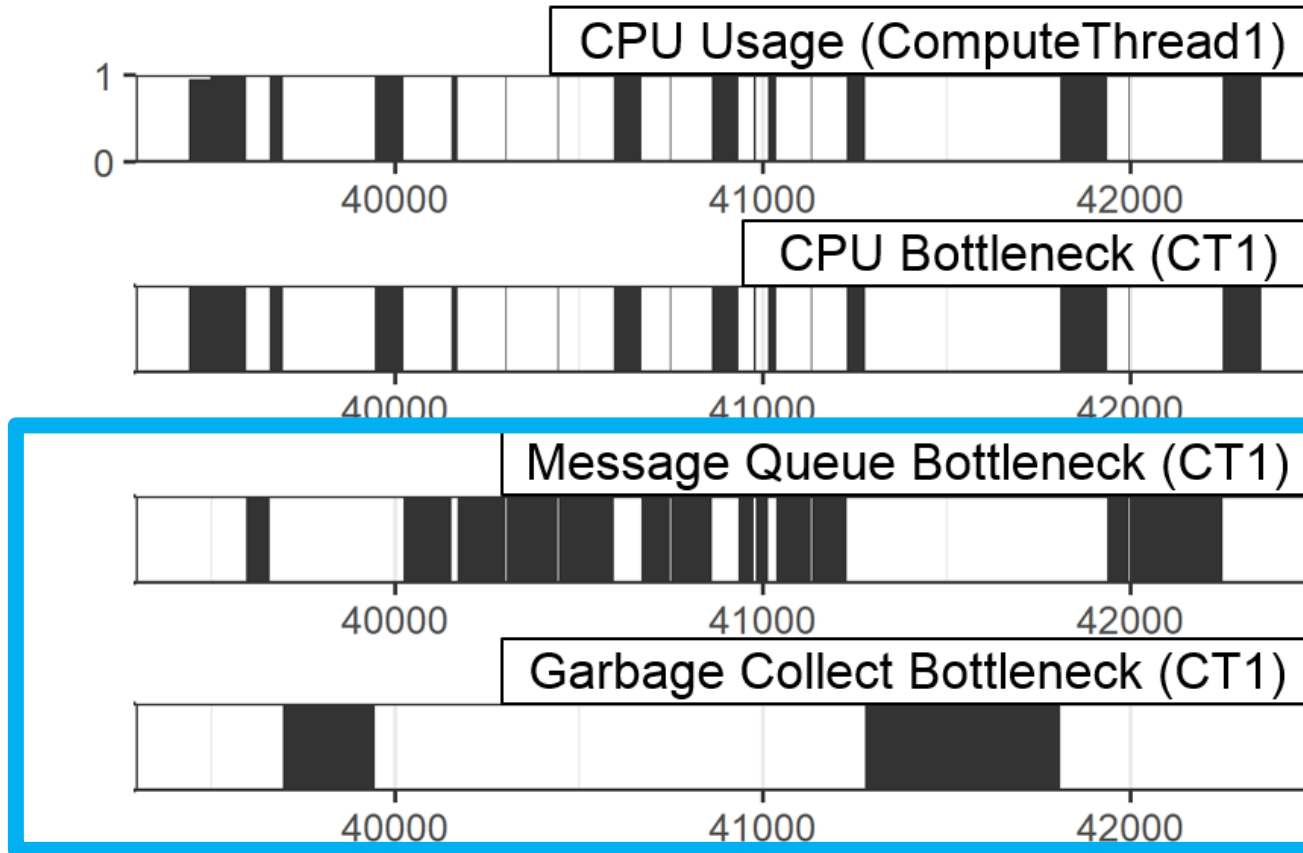
Preliminary Result: Analysing a Giraph Job



Preliminary Result: Analysing a Giraph Job



Preliminary Result: Analysing a Giraph Job



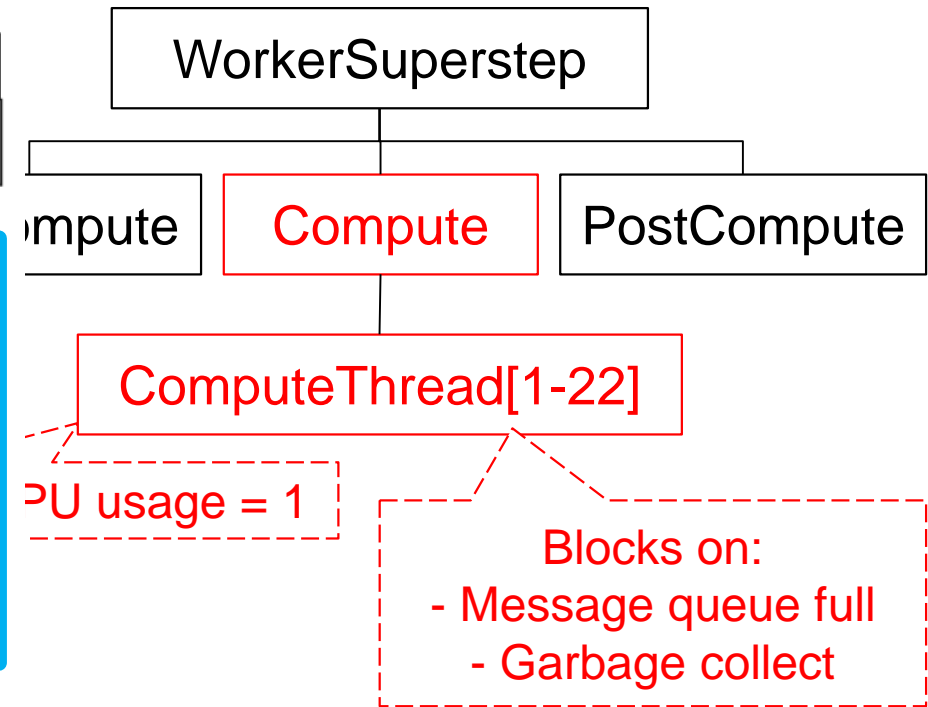
Grade10 : Help users understand the performance of graph-processing systems through automated analysis of performance data

Average time bottlenecked for Compute/ComputeThread:

- None: **0** ms (always bottlenecked)
- Message queue full: 1768 ms
- Garbage collect: 781 ms
- CPU: 748 ms

... So focus on reducing:

- Communication bottlenecks
- GC overheads (good luck!)



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[®]

References

1. Grade10:
(upcoming)
2. Granula:
[Ngai et al., GRADES'17] Ngai et al., Granula: Toward Fine-grained Performance Analysis of Large-scale Graph Processing Platforms. SIGMOD GRADES'17.
3. Graphalytics:
[Iosup et al., VLDB'16] Iosup et al. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. PVLDB 9(13): 1317-1328 (2016)
4. Performance study, pre-Graphalytics:
[Guo et al., CCGRID'15] Guo et al., An Empirical Performance Evaluation of GPU-Enabled Graph-Processing Systems. CCGRID 2015: 423-432
5. Performance study, pre-Graphalytics:
[Guo et al., IPDPS'14] Guo et al.: How Well Do Graph-Processing Platforms Perform? An Empirical Performance Evaluation and Analysis. IPDPS 2014: 395-404