LDBC Graphalytics v0.9
(graphalytics.org)

Wing Lung Ngai (Tech. lead 2016-2017),
Tim Hegeman, Stijn Heldens, Alexandru Iosup,

New members of the team : Alexandru Uta, Ahmed Musaafir,
Contributors: Arnau Prat-Pérez, Mihai Capotă, Petr Koupy, Yinglong Xia, Peter Boncz

@Large Research
Massivizing Computer Systems

TU Delft VU UNIVERSITY AMSTERDAM

Project History

- 2013 - Performance studies on distributed graph processing
- 2014 - IPDPS article
- 2015 - Performance studies on GPU-based graph processing
- 2015 - First prototype of Graphalytics (v0.2)
- 2016 - VLDB article (**v0.3**)
- 2017 (now) - The benchmark is ready! (**v0.9)**
- end 2017 -> Global competition (**v1.0**)

Benchmark Specification

(LDBC Graphalytics 0.9)

http://github.com/ldbc/ldbc_graphalytics_docs

Benchmark Specification

System-under-test (platform @ environment)

Platform (software):

- large diversity in architecture design, languages, paradigms

Totem

OpenG

PGX

Powergraph

GraphMat

Environment (hardware):

- community / high-end, many-cores / multi-cores

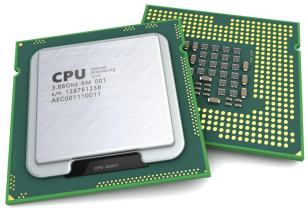

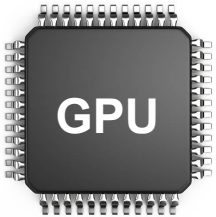

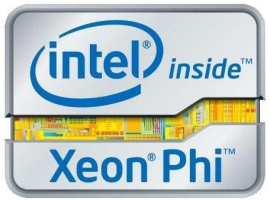

Benchmark Specification

Graph-processing workload

Label	Scale	Size
S	7.5-7.9	~50M
M	8.0-8.4	~160M
L	8.5-8.9	~500M
XL	9.0-9.4	~1.6B

Target-scale

For each *target-scale*

- 6 algorithms:
  - BFS, WCC, PR, CDLP, LCC, SSSP
- 5 datasets
  - both real-world and synthetic
- 5x repetitions
- In total, 150 benchmark runs.

A Renewable Process

Benchmark Specification

Graph-processing job & Performance metrics

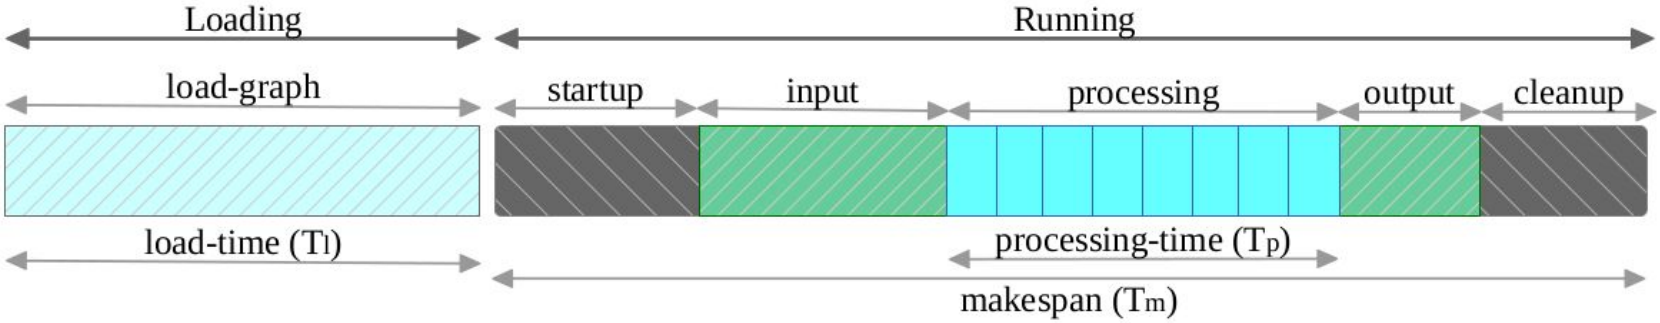

- Loading Time [seconds]
- Makespan [seconds]
- Processing Time [seconds]
- EVPS (edge and vertices per seconds) [unit]
- PPP (Price-per-performance) [dollar / unit]
- EPP (Energy-per-performance) [watts / unit] (v2.0)

User Participation

(LDBC Graphalytics 0.9)

http://graphalytics.org

User Participation

How to run the Graphalytics benchmark?

Our website
graphalytics.org

- documentation
- platform drivers (5+)
- benchmark resources

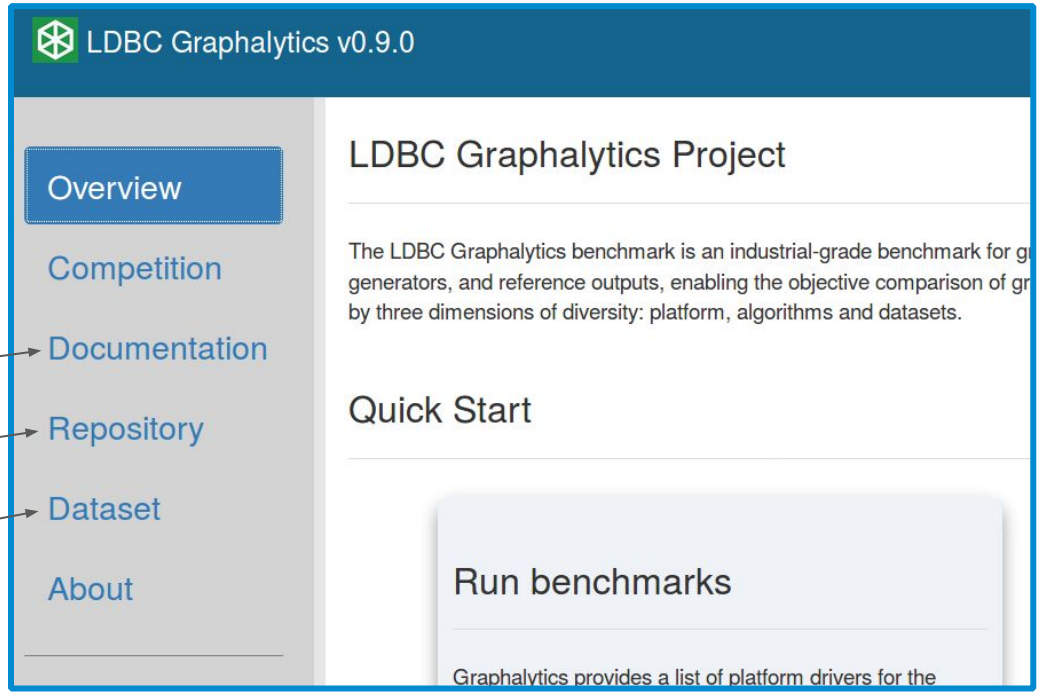

User Participation

How to add your platform driver?

Platform Variables
- platform_name="Xgraph"
- platform_acronym="xgraph"
- platform_version="1.0"
- developer_name="John Smith"

- Boilerplate code

- Platform interaction
- Algorithm implementation
- Code optimization

```
mvn archetype:generate -B \
 -DarchetypeGroupId=science.atlarge.graphalytics \
 -DarchetypeArtifactId=graphalytics-platforms-default-archetype \
 -DarchetypeVersion=0.9.0 \
 -DgroupId=science.atlarge.graphalytics \
 -Dpackage=science.atlarge \
 -DartifactId="graphalytics-platforms-${platform_acronym}"
 -Dversion=0.1-SNAPSHOT \
 -Dplatform-name="${platform_name}" \
 -Dplatform-acronym="${platform_acronym}" \
 -Dplatform-version="${platform_version}" \
 -Ddeveloper-name="${developer_name}"
```

Xgraph platform driver
- 2042 line of code
- 23 files

User Participation

How to view/submit benchmark results?

Json Result Format

```
"id": "b634810",
"specification": "0.9.0",
"system": {
      "platform": {...}, "environment": {...},
      "pricing": "..."
},
"benchmark": {
      "type": "standard",
      "target_scale": "L",
      ...
      }
},
"result": {
     ……
}
```

Full Disclosure Report

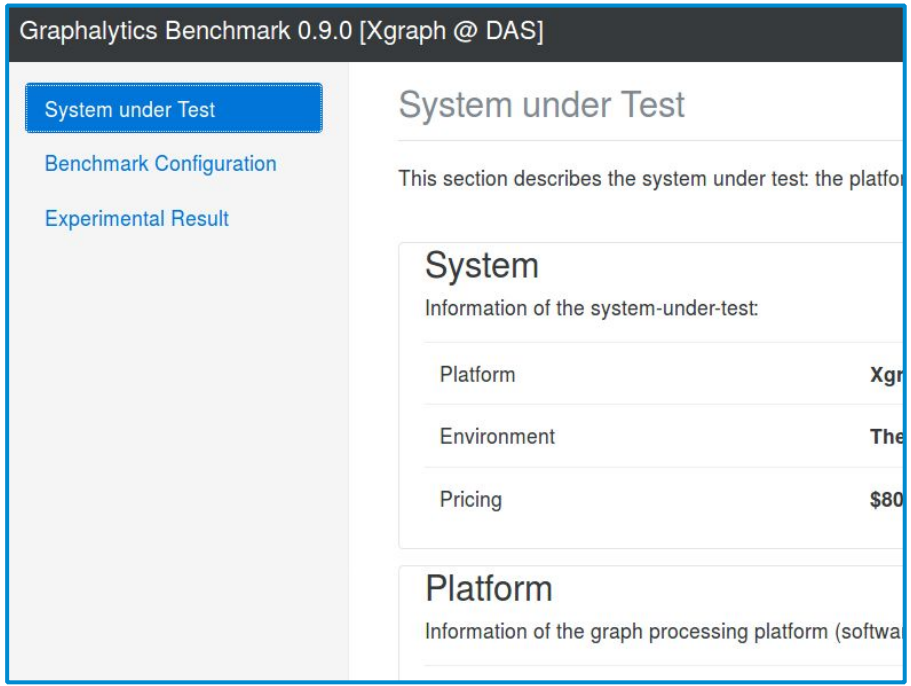

Global Competition
(LDBC Graphalytics v1.0)

Coming soon (ETA 1 - 3 months)

Global Competition

- LDBC global competition
  - coherent to the LDBC guideline (TPC pricing model)
  - ranking method: single value-of-merit
  - e.g. price-per-performance (PPP) score for targe-scale L

- Graphalytics global competition
  - broader participation (prototype to production)
  - ranking method: pair-wise comparison
  - e.g. diverse set of performance metrics, target-scales...

LDBC Graphalytics v0.9
(graphalytics.org)

Looking forward to your participation!