

Benchmarking platforms for large-scale graph processing and RDF data management

Bibliography

Main references:

1. Yong Guo, Marcin Biczak, Ana Lucia Varbanescu, Alexandru Iosup, Claudio Martella, Theodore L. Willke: *How Well Do Graph-Processing Platforms Perform? An Empirical Performance Evaluation and Analysis*.
In IPDPS 2014, Phoenix, Arizona, USA, May 2014.
2. Alexandru Iosup, Ana Lucia Varbanescu, Mihai Capota, Tim Hegeman, Yong Guo, W.-L. Ngai, Merijn Verstraaten: *Towards Benchmarking IaaS and PaaS Clouds for Graph Analytics*.
In WBDB 2014, Potsdam, Germany, Aug 2014.
3. Alexandru Iosup, Radu Prodan, and Dick Epema: *IaaS Cloud Benchmarking: Approaches, Challenges, and Experience*.
In MTAGS 2012 (with SC'12), Salt Lake City, Utah, USA, Nov 2012.
4. Ana Lucia Varbanescu, Merijn Verstraaten, Ate Penders, Alexandru Iosup, Henk Sips, and Cees de Laat: *Can Portability Improve Performance? An Empirical Study of Parallel Graph Analytics*.
In ICPE 2015, Austin, Texas, USA, Jan 2015.
5. Yong Guo, Ana Lucia Varbanescu, Alexandru Iosup and Dick Epema: *An Empirical Performance Evaluation of GPU-Enabled Graph-Processing Systems*.
In CCGrid 2015, May 2015.
6. Pham, M. D., Boncz, P., & Erling, O. : *S3g2: A scalable structure-correlated social graph generator*.
In Selected Topics in Performance Evaluation and Benchmarking (pp. 156-172), 2012.
7. Erling, O., Averbuch, A., Larriba-Pey, J., Chafi, H., Gubichev, A., Prat, A.,... & Boncz, P. (2015, May). *The LDBC Social Network Benchmark: Interactive Workload*.
In ACM SIGMOD, May 2015.
8. Capota, M., Hegeman, T., Iosup, A., Prat-Pérez, A., Erling, O., & Boncz, P. *Graphalytics: A Big Data Benchmark for Graph-Processing Platforms*.
In GRADES, with ACM SIGMOD, May 2015.

Recommended reading about the systems included in the tutorial:

6. A. Gharaibeh, E. Santos-Neto, L. B. Costa, and M. Ripeanu: **A Yoke of Oxen and a Thousand Chickens for Heavy Lifting Graph Processing**.
In PACT 2012, San Francisco, California, USA.

7. A. Gharaibeh, E. Santos-Neto, L. B. Costa, and M. Ripeanu: **Efficient Large-Scale Graph Processing on Hybrid CPU and GPU Systems.**
In TOPC, 2013.
8. J. Zhong and B. He: **Medusa: Simplified Graph Processing on GPUs.**
In TPDS, 2013.
9. Z. Fu, M. Personick, and B. Thompson: **MapGraph: A High Level API for Fast Development of High Performance Graph Analytics on GPUs.**
In GRADES, 2014.
10. G. Malewicz, M. H. Austern, A. J. Bik, J. C. Dehnert, I. Horn, N. Leiser, and G. Czajkowski: **Pregel: A System for Large-scale Graph Processing.**
In SIGMOD, 2010.
11. Y. Low, D. Bickson, J. Gonzalez, C. Guestrin, A. Kyrola, and J. M. Hellerstein: **Distributed GraphLab: A Framework for Machine Learning and Data Mining in the Cloud.** In VLDB, 2012.
12. J. E. Gonzalez, R. S. Xin, A. Dave, D. Crankshaw, M. J. Franklin, and I. Stoica: **GraphX: Graph Processing in a Distributed Dataflow Framework.**
In OSDI, 2014.
13. Alexander Alexandrov , Rico Bergmann , Stephan Ewen , Johann-Christoph Freytag , Fabian Hueske , Arvid Heise , Odej Kao , Marcus Leich , Ulf Leser , Volker Markl , Felix Naumann , Mathias Peters , Astrid Rheinländer , Matthias J. Sax , Sebastian Schelter , Mareike Höger , Kostas Tzoumas , Daniel Warneke: **The Stratosphere platform for Big Data Analytics.**
In VLDB Journal, 2014.

Code and documentation

- Graphalytics

<https://github.com/tudelft-atlarge/graphalytics/>

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

- The latest LDBC SNB benchmark specification

https://github.com/ldbc/ldbc_snb_docs/blob/master/LDBC_SNB_v0.2.2.pdf

- Other systems:

Totem: <http://netsyslab.ece.ubc.ca/wiki/index.php/Totem>

MapGraph: <http://mapgraph.io>

Medusa: <https://code.google.com/p/medusa-gpu/>

YARN: <http://hadoop.apache.org/docs/current/hadoop-yarn/hadoop-yarn-site/YARN.html>

Neo4J: <http://www.neo4j.org/>

Stratosphere: <http://stratosphere.eu/>

<http://flink.apache.org>

Giraph: <http://giraph.apache.org>

GraphLab/Dato: <https://dato.com/>

GraphX: <http://spark.apache.org/graphx/>