

Towards Benchmarking Graph-Processing Platforms

Yong Guo, Marcin Biczak, Ana Lucia Varbanescu,
Alexandru Iosup, Claudio Martella, and Theodore L. Willke

Yong.Guo@tudelft.nl

The Data Deluge



Too Many Untested Platforms



Benchmarking Suite

Performance metrics

- Raw processing power
- Resource utilization
- Scalability
- Overhead

Graph	#V	#E	Size
Amazon	262.1 K	1.2 M	18 MB
WikiTalk	2.4 M	5.0 M	87 MB
KGS	293.3K	16.6 M	210 MB
Citation	3.8 M	16.5 M	297 MB
DotaLeague	61.2 K	50.9 M	655 MB
Synth	2.4 M	64.2 M	964 MB
Friendster	65.6 M	1.8 B	31 GB

Algorithms


- **STATS**: single step, overview
- **BFS**: iterative, building block
- **CONN**: iterative, building block
- **CD**: iterative, social network
- **EVO**: multi-level iterative, prediction

Results & Analysis




DAS4: a multi-cluster Dutch grid/cloud


- Intel Xeon 2.4 GHz CPU (12 MB cache)
- Memory 24 GB
- 10 Gbit/s Infiniband and 1 Gbit/s Ethernet
- 20-50 working nodes used in our experiments



The execution time of the BFS algorithm



The horizontal scalability of processing Friendster



The execution time breakdown for BFS



Key findings

- Performance is a function of (Dataset, Algorithm, Platform, Deployment)
- Cannot tell yet which of (Dataset, Algorithm, Platform) is the most important
- Platforms drawbacks (e.g., long execution time, crashes, tuning, etc.)
- Some platforms can scale up reasonably with cluster size or number of cores



Is this a benchmarking suite?

Ask me and let's **discuss**



Join us within the SPEC RG Cloud Working Group

<http://research.spec.org/working-groups/rg-cloud-working-group.html>

Y. Guo, M. Biczak, A. L. Varbanescu, A. Iosup, C. Martella, and T. L. Willke, "How Well do Graph-Processing Platforms Perform? An Empirical Performance Evaluation and Analysis: Extended Report," TU Delft, PDS-2013-004, 2013, <http://www.pds.ewi.tudelft.nl/researchpublications/technical-reports/2013/>.