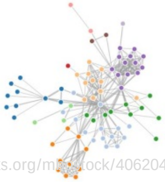
http://bl.ocks.org/mbostock/4062045

GRAPHALYTICS
A Big Data Benchmark for Graph-Processing Platforms

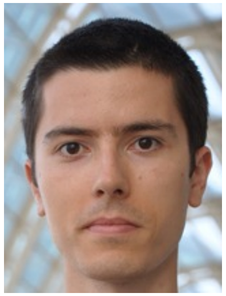
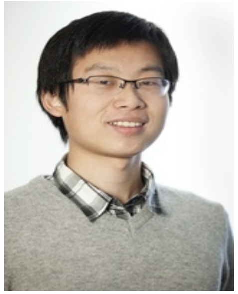


Mihai Capotã, Yong Guo, Ana Lucia Varbanescu,

Tim Hegeman,
Jorai Rijsdijk,

GRAPHALYTICS was made possible by a generous contribution from Oracle.

Alexandru Iosup,

Jose Larriba Pey, Arnau Prat, Peter Boncz, Hassan Chafi

(TU) Delft – the Netherlands – Europe

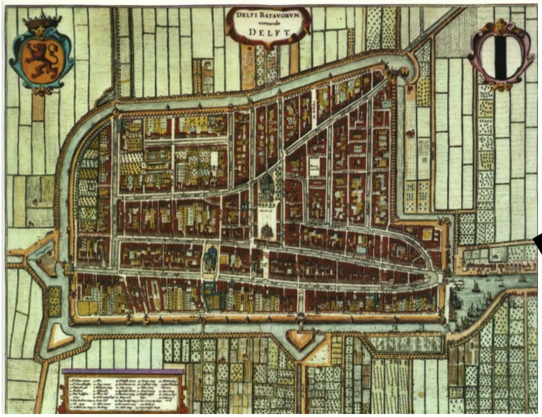

founded 13th century
pop: 100,000

founded 1842
pop: 13,000

Delft

pop: 16.5 M

Barcelona

The Parallel and Distributed Systems Group at TU Delft

VENI

Alexandru Iosup

Grids/Clouds
P2P systems
Big Data/graphs
Online gaming

Dick Epema

Grids/Clouds
P2P systems
Video-on-demand
e-Science

VENI

Ana Lucia Varbanescu (now UvA)

HPC systems
Multi-cores
Big Data/graphs

Henk Sips

HPC systems
Multi-cores
P2P systems

VENI

Johan Pouwelse

P2P systems
File-sharing
Video-on-demand

Home page

- www.pds.ewi.tudelft.nl

Publications

- see PDS publication database at publications.st.ewi.tudelft.nl

tribler

The Koala Grid Scheduler

Open TTD @large

GrenchMark

The Grid Workloads Archive

Winners IEEE TCSC Scale Challenge 2014

TUDelft
Delft University of Technology

Graphs at the Core of Our Society: The LinkedIn Example

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

100M Mar 2011, 69M May 2010

Sources: Vincenzo Cosenza, The State of LinkedIn, http://vincos.it/the-state-of-linkedin/
via Christopher Penn, http://www.shiftcomm.com/2014/02/state-linkedin-social-media-dark-horse/

Graphs at the Core of Our Society: The LinkedIn Example

Sources: Vincenzo Cosenza, The State of LinkedIn, http://vincos.it/the-state-of-linkedin/
via Christopher Penn, http://www.shiftcomm.com/2014/02/state-linkedin-social-media-dark-horse/

The data deluge: large-scale graphs

LinkedIn

300M users

??? edges

YAHOO!

friendster

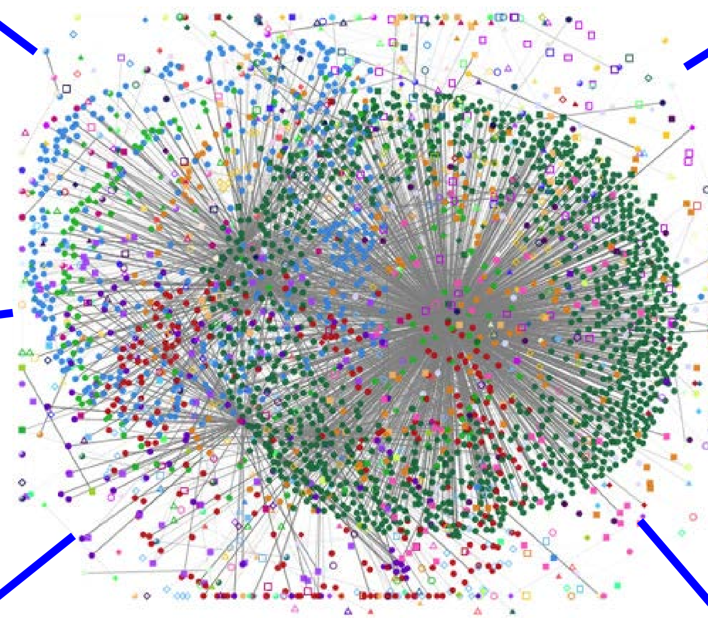

Twitter

270M MAU
200+ avg followers

>54B edges

1.2B MAU 0.8B DAU
200+ avg followers

>240B edges

XFIRE

The data deluge: large-scale graphs

Oracle 1.2M followers,
132k employees

company/day:
40-60 posts, 500-700 comments

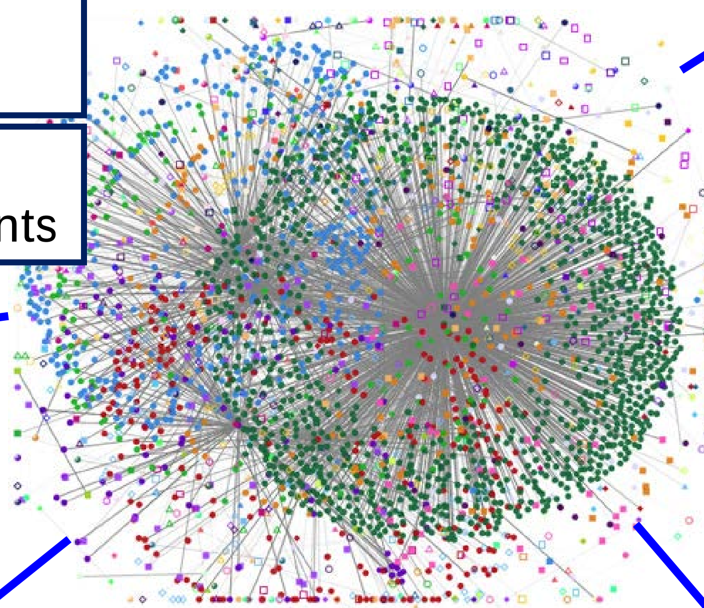

270M MAU
200+ avg followers

>54B edges

1.2B MAU 0.8B DAU
200+ avg followers

>240B edges

The data deluge: large-scale graphs

Oracle 1.2M followers

270M MAU

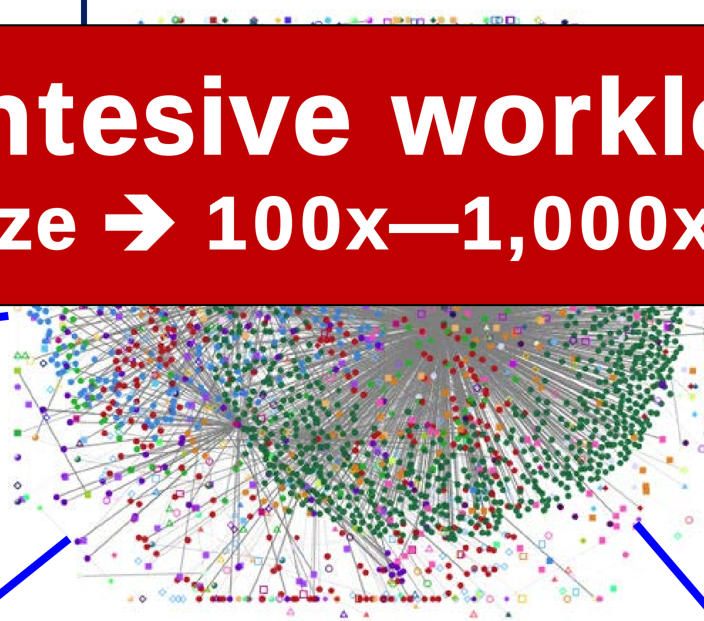

Data-intesive workload

10x graph size → 100x—1,000x slower

1.2B MAU 0.8B DAU
200+ avg followers

>240B edges

The data deluge: large-scale graphs

Oracle 1.2M followers

270M MAU

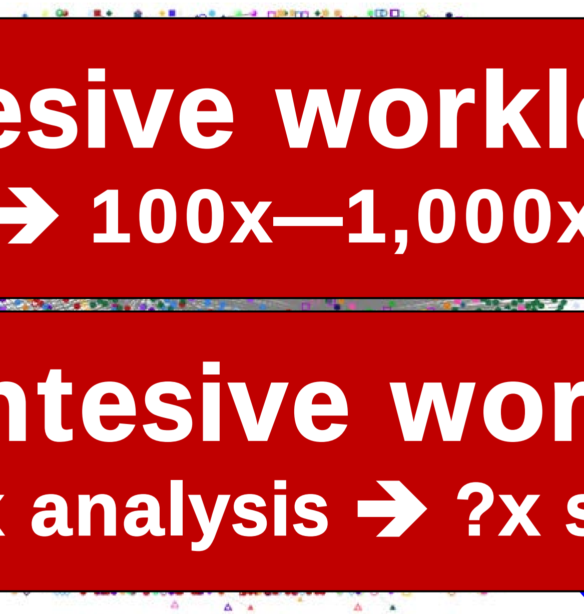

Data-intesive workload

10x graph size → 100x—1,000x slower

Compute-intesive workload

more complex analysis → ?x slower

>240B edges

The data deluge: large-scale graphs

Oracle 1.2M followers

270M MAU

Data-intesive workload

10x graph size → 100x—1,000x slower

Compute-intesive workload

more complex analysis → ?x slower

Dataset-dependent workload

unfriendly graphs → ??x slower

Graphs at the Core of Our Society: The LinkedIn Example

The State of **LinkedIn**

3-4 new users every second

but fewer visitors (and page views)

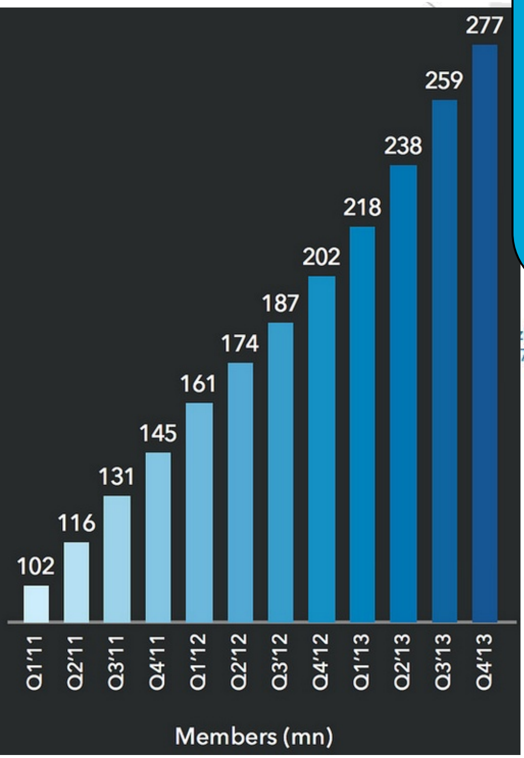

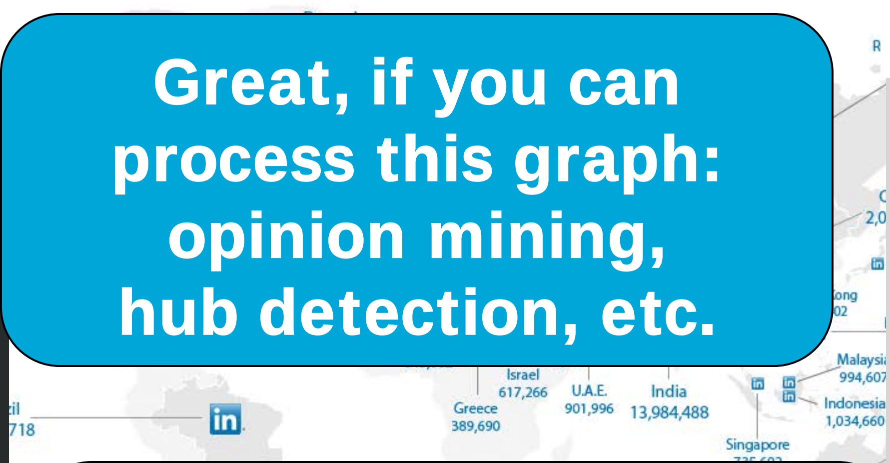

Great, if you can process this graph: opinion mining, hub detection, etc.

100+ million questions of customer retention, of (lost) customer influence, of ...

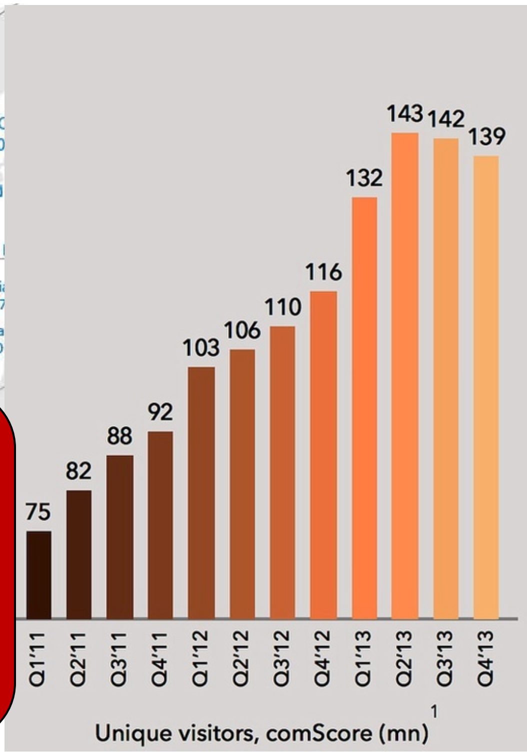

Sources: Vincenzo Cosenza, The State of LinkedIn, http://vincos.it/the-state-of-linkedin/
via Christopher Penn, http://www.shiftcomm.com/2014/02/state-linkedin-social-media-dark-horse/

TU Delft
Delft University of Technology

Graphs at the Core of Our Society: The LinkedIn Example

The State of LinkedIn

3-4 new users every second

but fewer visitors (and page views)

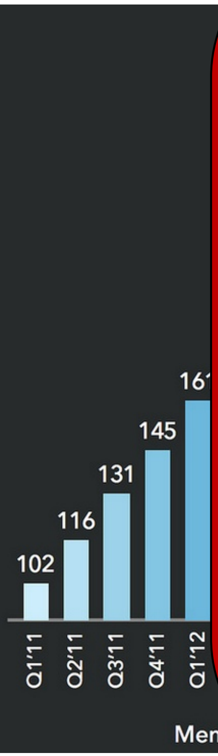

Periodic and/or continuous analytics at full scale

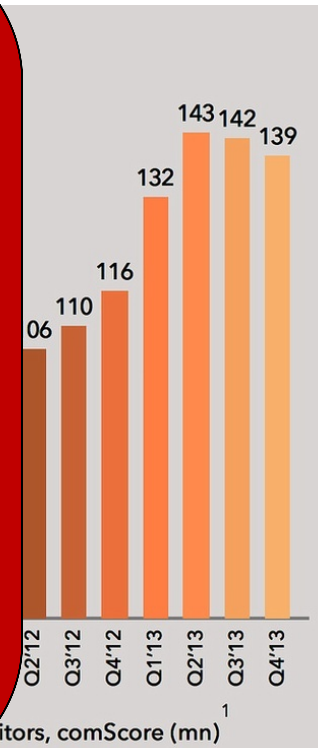

Sources: Vincenzo Cosenza, The State of LinkedIn, http://vincos.it/the-state-of-linkedin/
via Christopher Penn, http://www.shiftcomm.com/2014/02/state-linkedin-social-media-dark-horse/

LinkedIn Maps **Ana Lucia Varbanescu's Professional Network**
as of November 28, 2013

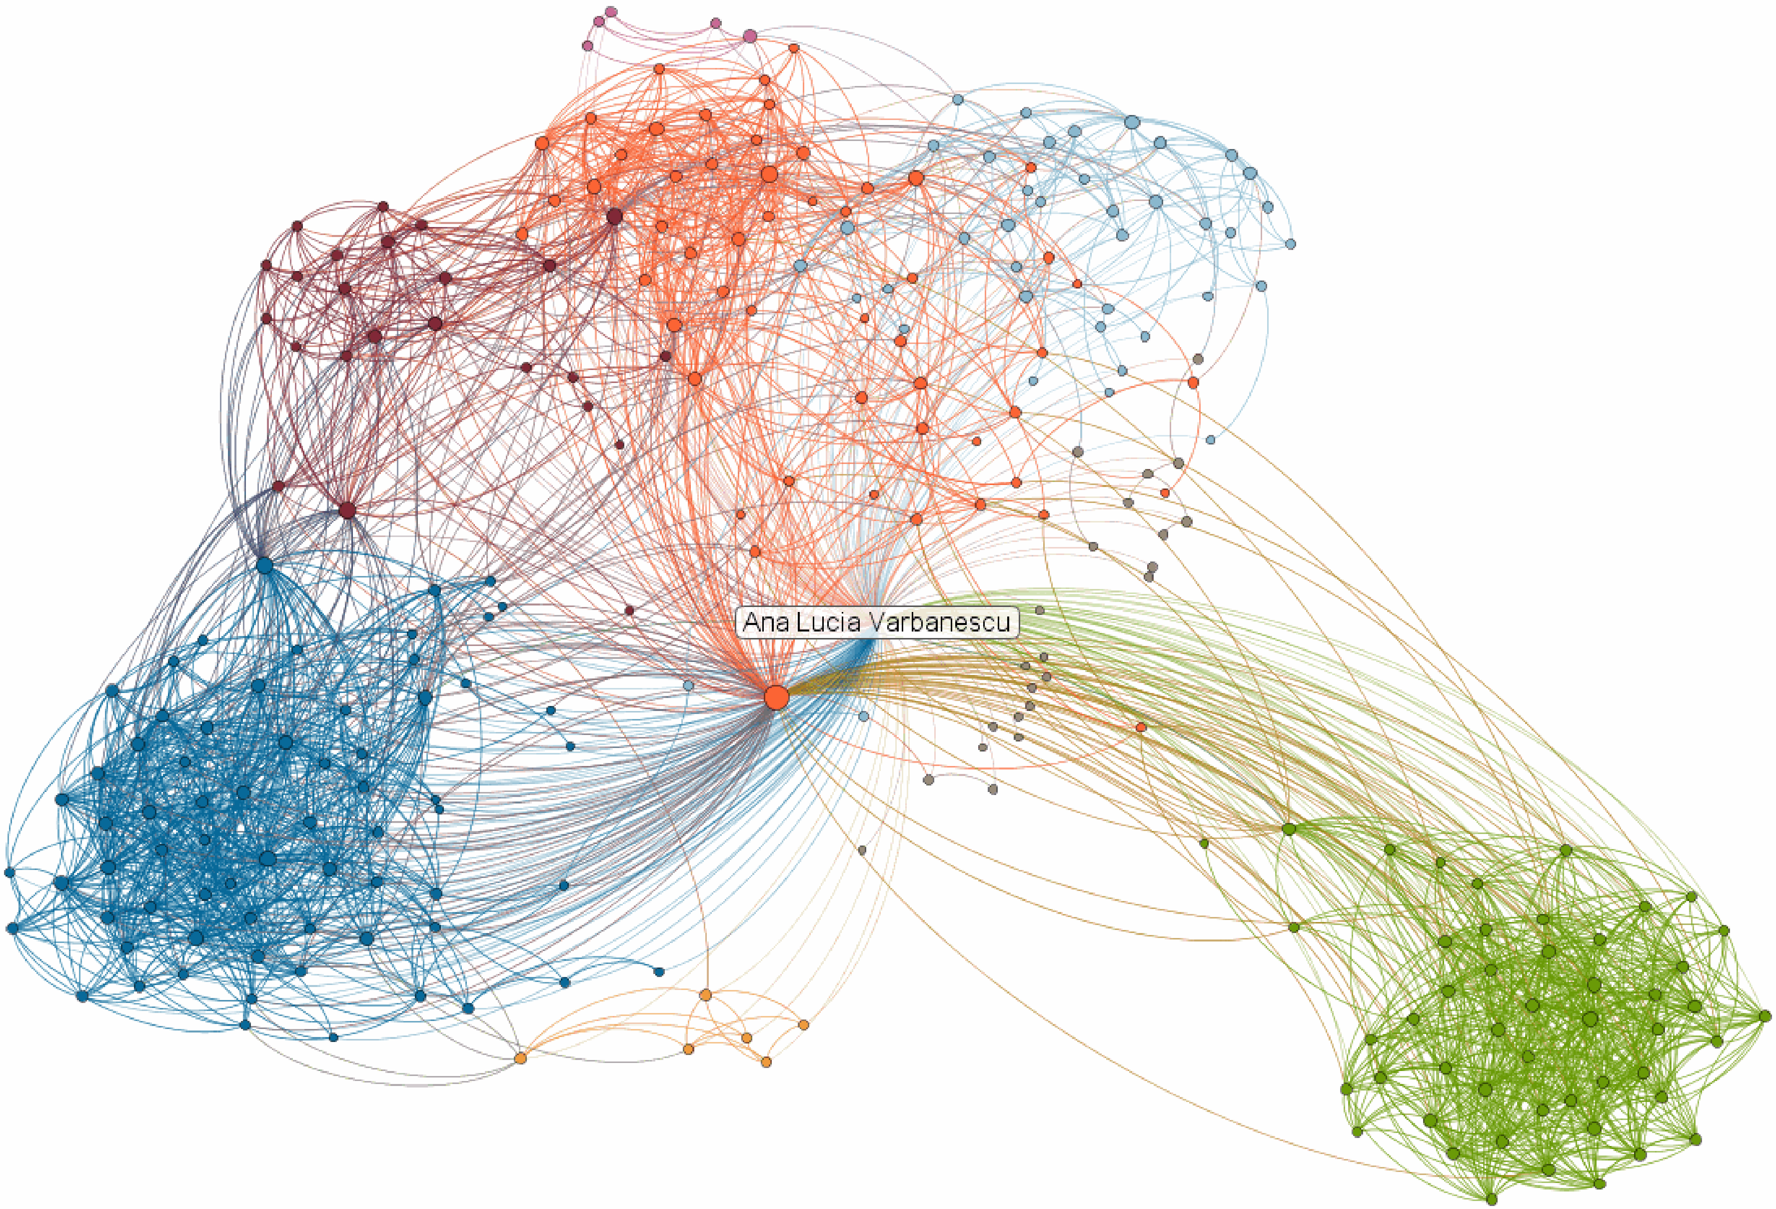

The “sorry, but...” moment

Supporting multiple users

10x number of users ➔ ????x slower

Graph Processing @large

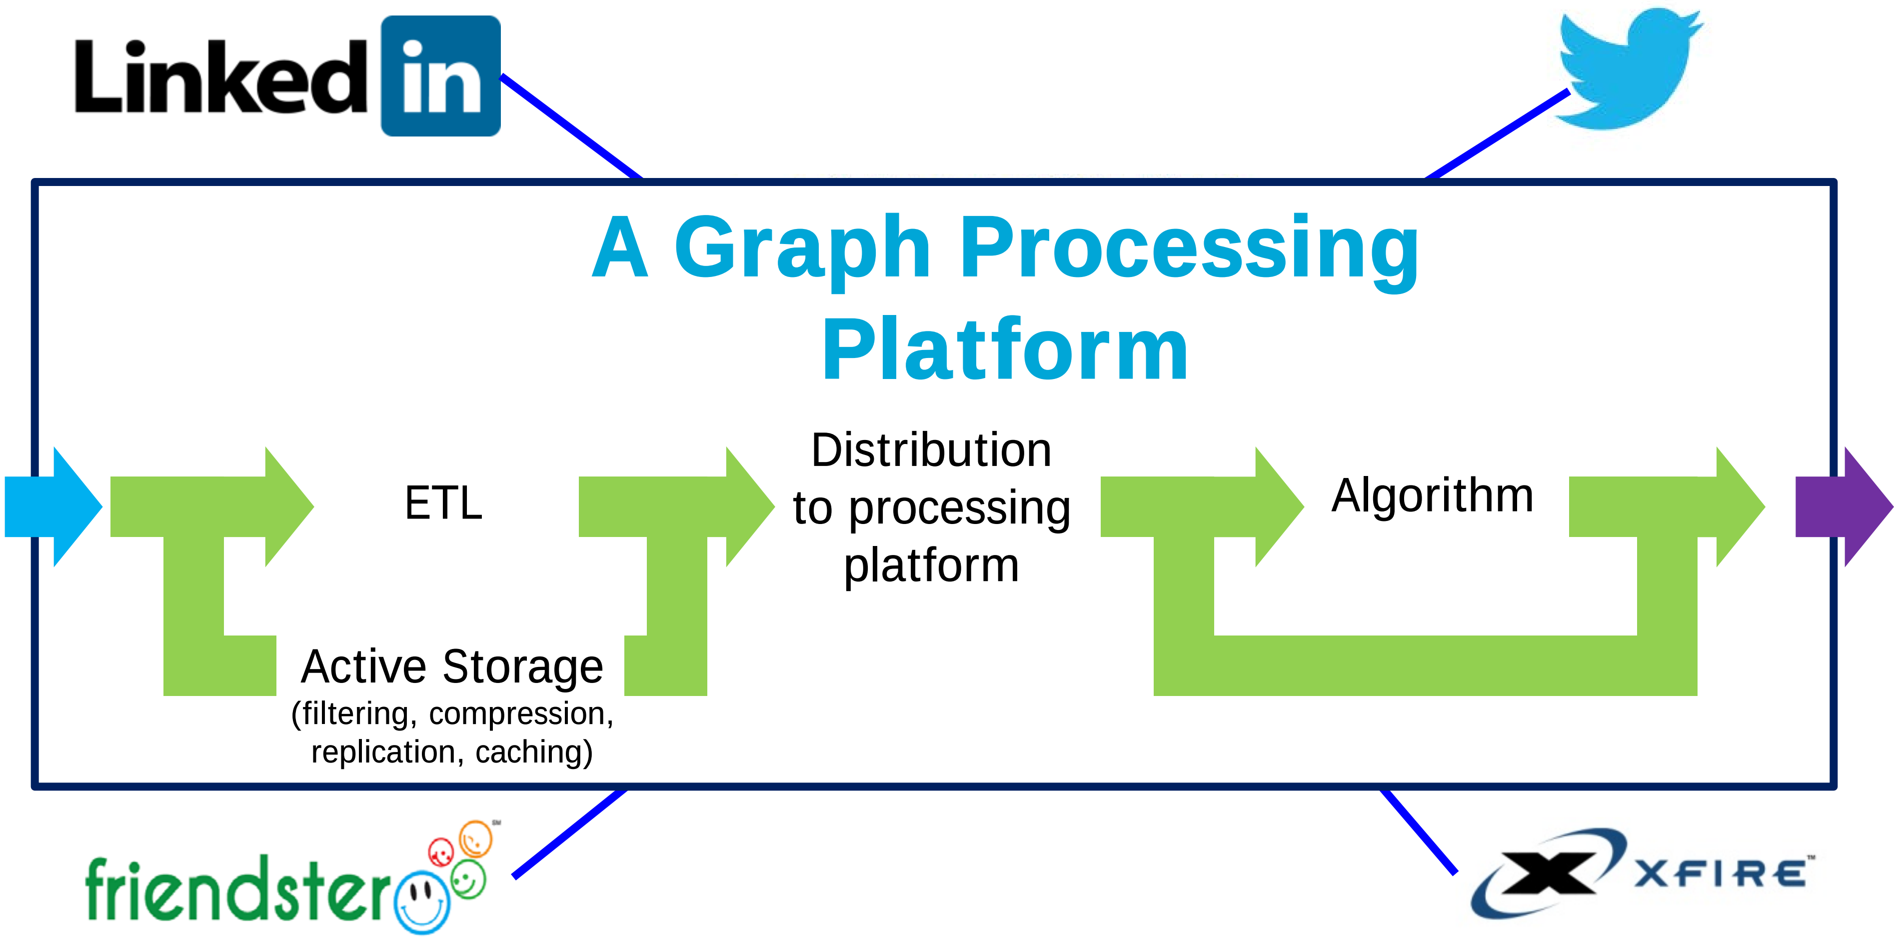

Interactive processing not considered in this presentation.
Streaming not considered in this presentation.

Graph Processing @large

A Graph Processing Platform

Ideally, N cores/disks → Nx faster

replication, caching)

Distribution to processing platform

Ideally, N cores/disks → Nx faster

Interactive processing not considered in this presentation.

Streaming not considered in this presentation.

Graph-Processing Platforms

- Platform: the combined hardware, software, and programming system that is being used to complete a graph processing task

Which to choose?
What to tune?

What is the performance of graph-processing platforms?

Metrics Diversity

Graph Diversity

Algorithm Diversity

Graphalytics = comprehensive benchmarking suite for graph processing across all platforms

Graphalytics = A Challenging Benchmarking Process

- Methodological challenges
  - Challenge 1. Evaluation process
  - Challenge 2. Selection and design of performance metrics
  - Challenge 3. Dataset selection and analysis of coverage
  - Challenge 4. Algorithm selection and analysis of coverage

- Practical challenges
  - Challenge 5. Scalability of evaluation, selection processes
  - Challenge 6. Portability
  - Challenge 7. Result reporting

Y. Guo, A. L. Varbanescu, A. Iosup, C. Martella, T. L. Willke: Benchmarking graph-processing platforms: a vision. ICPE 2014: 289-292

Graphalytics = Many Classes of Algorithms

- Literature survey of of metrics, datasets, and algorithms
  - 10 top research conferences: SIGMOD, VLDB, HPDC ...
  - Key word: graph processing, social network
  - 2009–2013, 124 articles

Class	Examples	%
Graph Statistics	Diameter, PageRank	16.1
Graph Traversal	BFS, SSSP, DFS	46.3
Connected Component	Reachability, BiCC	13.4
Community Detection	Clustering, Nearest Neighbor	5.4
Graph Evolution	Forest Fire Model, PAM	4.0
Other	Future work	

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, IPDPS'14.

Graphalytics = Real & Synthetic Datasets

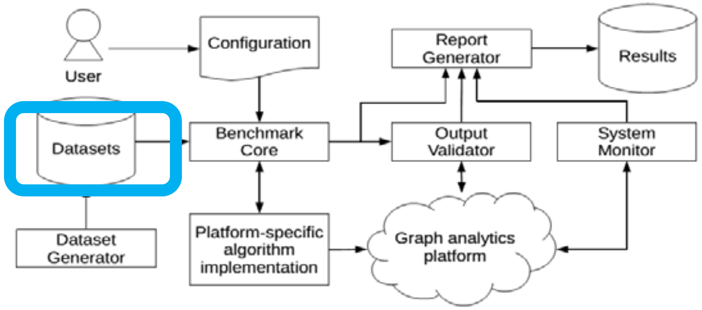

	Graphs	#V	#E	d	D̄	Directivity
G1	Amazon	262,111	1,234,877	1.8	4.7	directed
G2	WikiTalk	2,388,953	5,018,445	0.1	2.1	directed
G3	KGS					undirected
G4	Citation					directed
G5	DotaLeague					undirected
G6	Synth					undirected
G7	Friendster	65,608,366	1,806,067,135	0.1	55.1	undirected

Interaction graphs
(possible work)

https://snap.stanford.edu/

LDBC
Social Network
Generator

The Game Trace Archive

http://gta.st.ewi.tudelft.nl/

Y. Guo and A. Iosup. The Game Trace Archive, NETGAMES 2012.

TU Delft
Delft University of Technology

Graphalytics = Advanced Harness

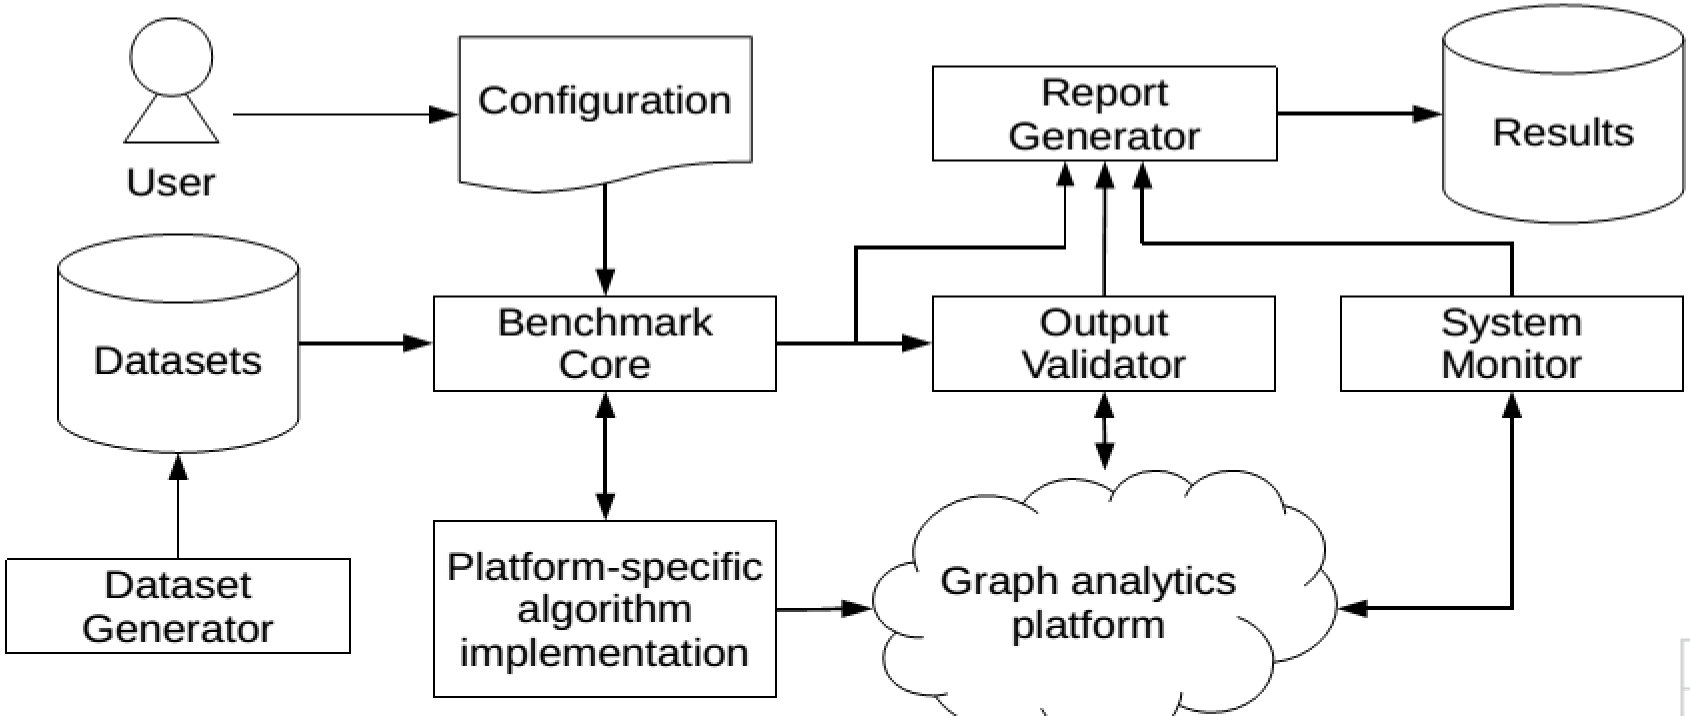

Cloud support technically feasible, methodologically difficult

Graphalytics = Enhanced LDBC Datagen

- A battery of graphs covering a rich set of configurations
- Datagen extensions to
  - More diverse degree distributions
  - Clustering coefficient and assortativity

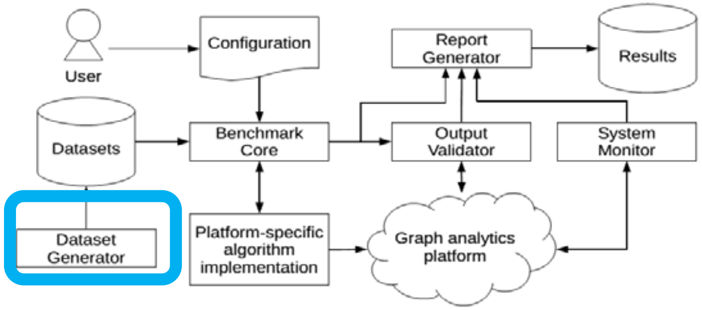

Ongoing work

LDBC D3.3.34 http://ldbcouncil.org/sites/default/files/LDBC_D3.3.34.pdf and Orri Erling et al. The LDBC Social Network Benchmark: Interactive Workload, SIGMOD'15

Graphalytics = Advanced Monitoring & Logging System

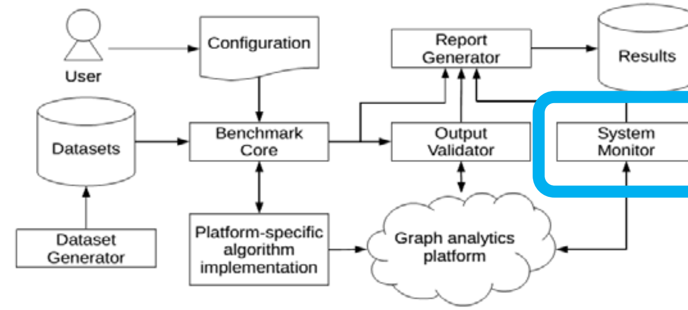

Diverse metrics: CPU, IOPS, Network, Memory use, …

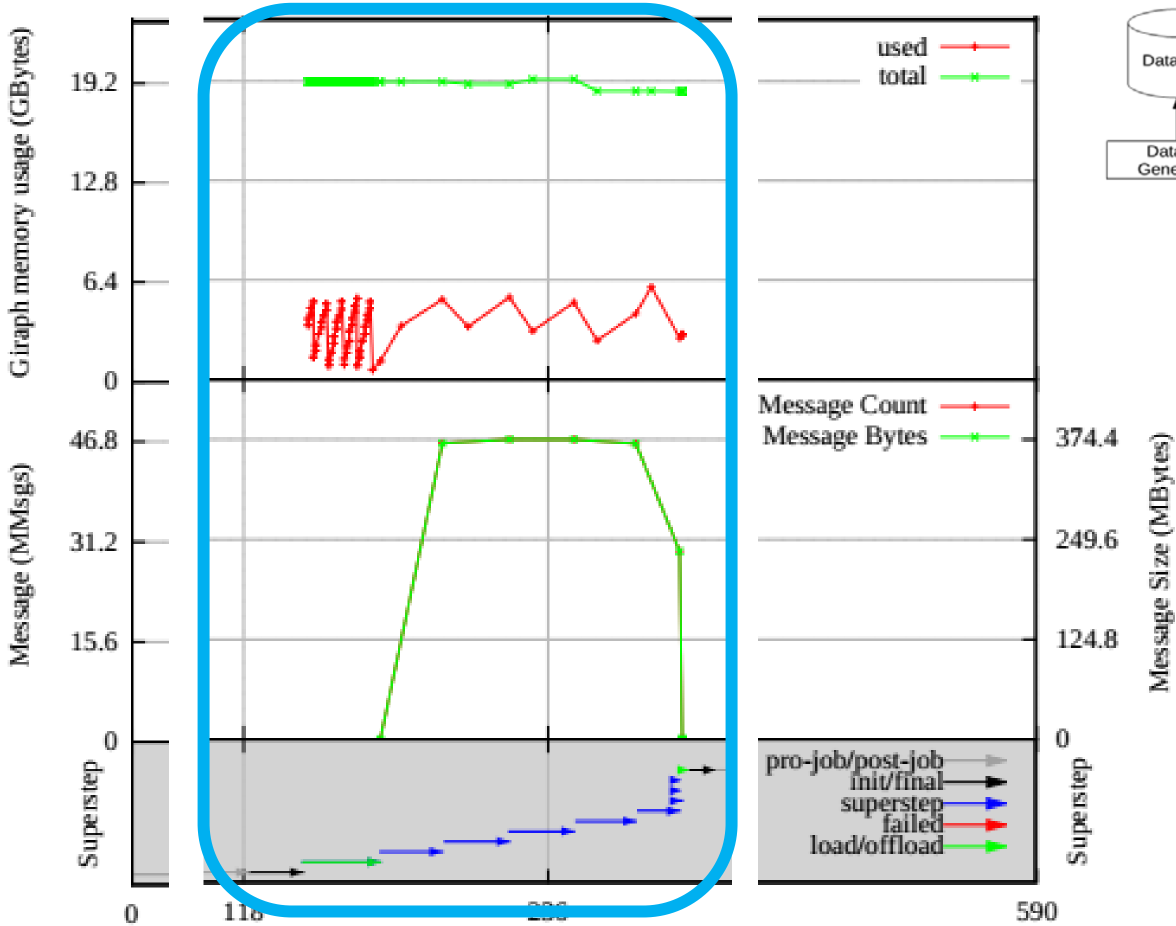

Time

- Automatic analysis matching the programming model

Ongoing work

A. Iosup et al., Towards Benchmarking IaaS and PaaS Clouds for Graph Analytics. WBDB 2014

TU Delft
Delft University of Technology

Graphalytics = Choke-Point Analysis

- Choke points are crucial technological challenges that platforms are struggling with

- Examples
  - Network traffic
  - Access locality
  - Skewed execution

- Challenge: Select benchmark workload based on real-world scenarios, but make sure they cover the important choke points

`near-future work`

Graphalytics = Advanced Software Engineering Process

https://github.com/mihaic/graphalytics/

- All significant modifications to Graphalytics are peer-reviewed by developers
  - Internal release to LDBC partners (Feb 2015)
  - Public release, announced first through LDBC (Apr 2015*)
- Jenkins continuous integration server
- SonarQube software quality analyzer

Graphalytics in Practice

6 real-world datasets +
2 synthetic generators

Data ingestion not included here!

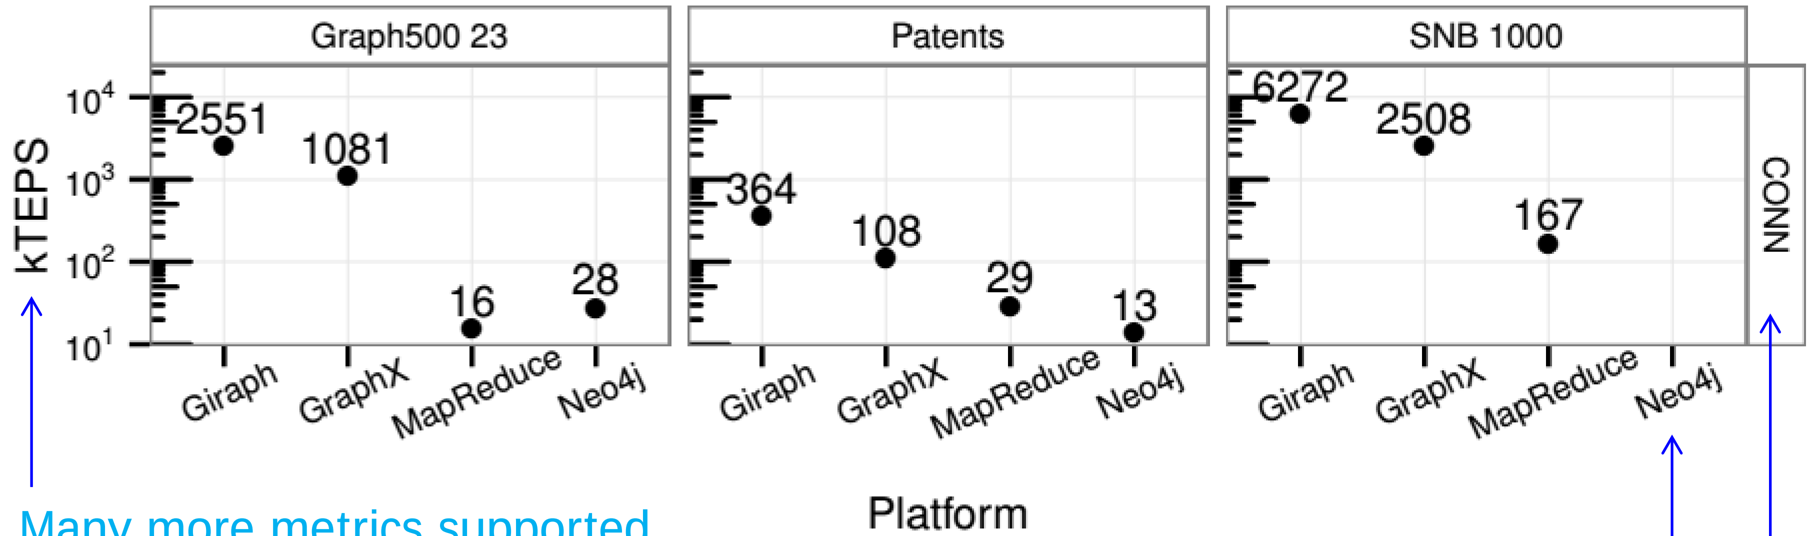

Many more metrics supported

10 platforms tested w prototype implementation

5 classes of algorithms

- Missing results = failures of the respective systems

Key Findings So Far

- Performance is function of (Dataset, Algorithm, Platform, Deployment)
  - Previous performance studies lead to tunnel vision
- Platforms have their specific drawbacks (crashes, long execution time, tuning, etc.)
  - Best-performing system depends on stakeholder needs
- Some platforms can scale up reasonably with cluster size (horizontally) or number of cores (vertically)
  - Strong vs weak scaling still a challenge—workload scaling tricky
  - Single-algorithm is not workflow/multi-tenancy

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,IPDPS’14.

Thank you for your attention!
Comments? Questions? Suggestions?

http://graphalytics.ewi.tudelft.nl
https://github.com/mihaic/graphalytics/

PELGA 2015, May 15
http://sites.google.com/site/pelga2015/

Alexandru Iosup
A.Iosup@tudelft.nl

GRAPHALYTICS was made possible by a generous contribution from Oracle.

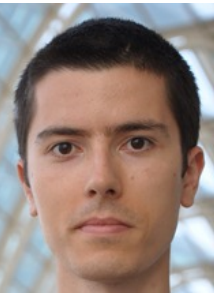
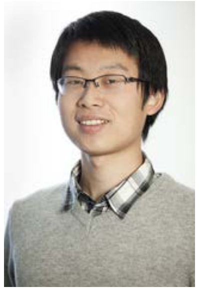
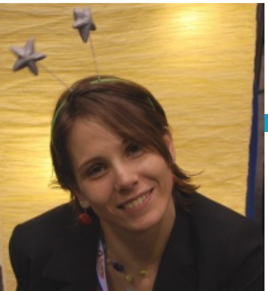

Contributors

A few extra slides

Discussion

- How much preprocessing should we allow in the ETL phase?
- How to choose a metric that captures the preprocessing?

http://graphalytics.ewi.tudelft.nl

Discussion

- How should we asses the correctness of algorithms that produce approximate results?
- Are sampling algorithms acceptable as trade-off time to benchmark vs benchmarking result?

http://graphalytics.ewi.tudelft.nl

Discussion

- How to setup the platforms? Should we allow algorithm-specific platform setups or should we require only one setup to be used for all algorithms?

http://graphalytics.ewi.tudelft.nl

Discussion

- Towards full use cases, full workflows, and inter-operation of big data processing systems
- How to benchmark the entire chain needed to produce useful results, perhaps even the human in the loop?

http://graphalytics.ewi.tudelft.nl

A. Iosup, T. Tannenbaum, M. Farrellee, D. H. J. Epema, M. Livny: Inter-operating grids through Delegated MatchMaking. Scientific Programming 16(2-3): 233-253 (2008)