LDBC The graph & RDF benchmark reference

# GRAPHALYTICS

## A Big Data Benchmark for Graph-Processing Platforms

**https://github.com/tudelft-atlarge/graphalytics/**

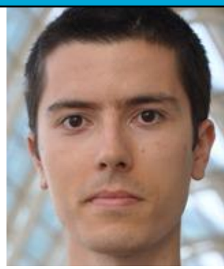

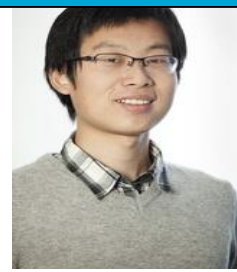

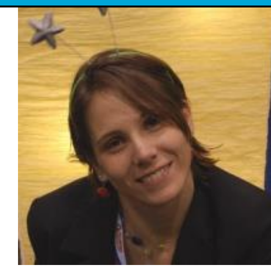


**Mihai Capotã, Yong Guo, Ana Lucia Varbanescu,**

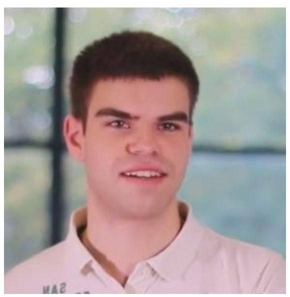

**Tim Hegeman,**

**Wing Lung Ngai,**

**Alexandru Iosup,**

UPC UPC CWI ORACLE

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



GRAPHALYTICS was made possible by a generous contribution from Oracle.

# (TU) Delft – the Netherlands – Europe

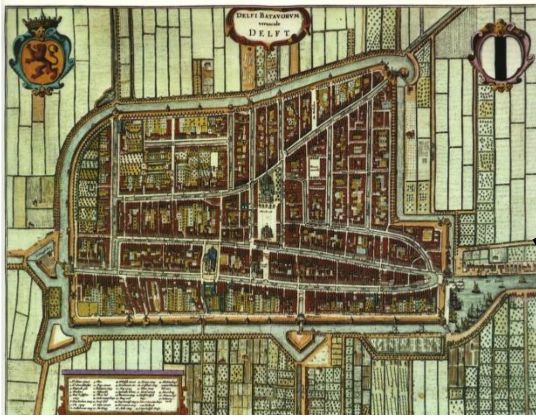


founded 13th century
pop: 100,000

founded 1842
pop: 15,000

Delft


pop: 16.5 M


Barcelona

# The Parallel and Distributed Systems Group at TU Delft


Alexandru Iosup

Grids/Clouds
P2P systems
Big Data/graphs
Online gaming

Dick Epema

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


Ana Lucia Varbanescu (now UvA)

HPC systems
Multi-cores
Big Data/graphs

Henk Sips

HPC systems
Multi-cores
P2P systems


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**

TU Delft — Delft University of Technology

# Graphs Are 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**

**Tens of "specialized LinkedIns": medical, mil, edu, gov, ...**



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

# LinkedIn’s Service/Ops Analytics

3-4 new users every second

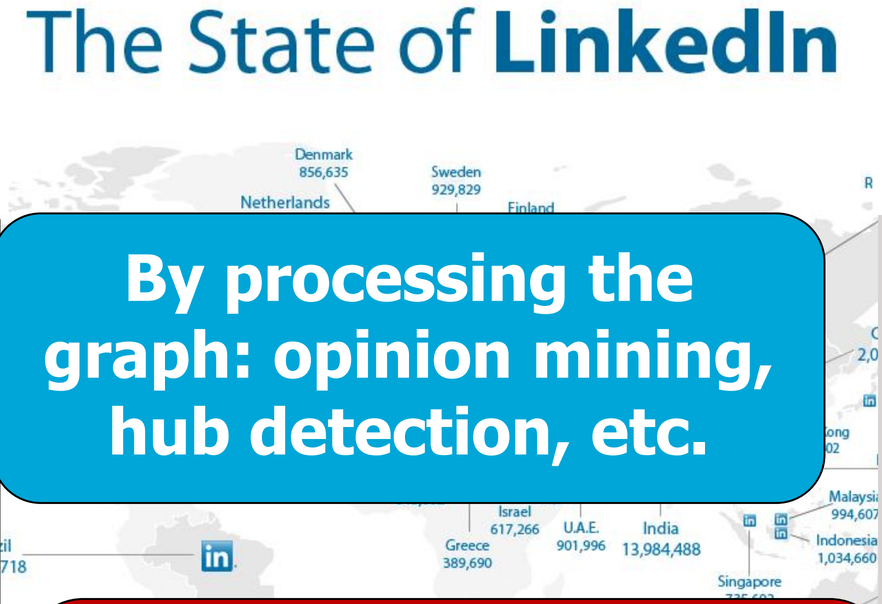


but fewer visitors (and page views)

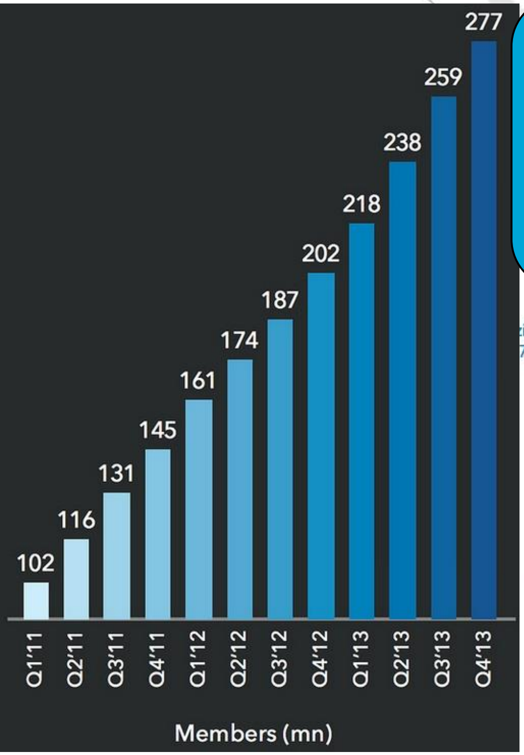


**By processing the graph: opinion mining, hub detection, etc.**

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

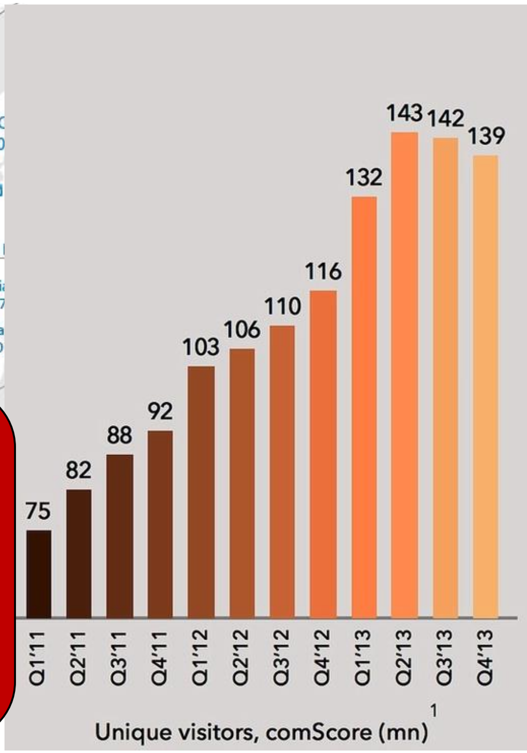




Sources: Vincenzo Cosenza, The State of LinkedIn, http://vincos.it/the-state-of-linkedin/
via Christopher Penn, http://www.shiftcomm.com/2014/02/state-linkedin-social-media-dark-horse/

# LinkedIn Analytics

## The State of **LinkedIn**

3-4 new users every second

but fewer visitors (and page views)

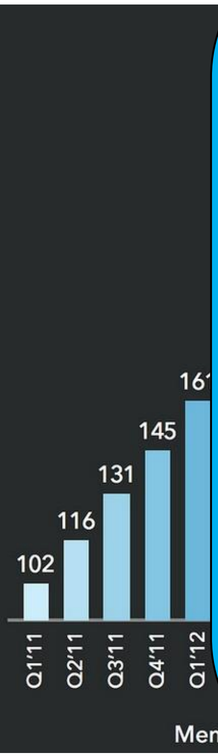


**Periodic and/or continuous analytics at full scale**

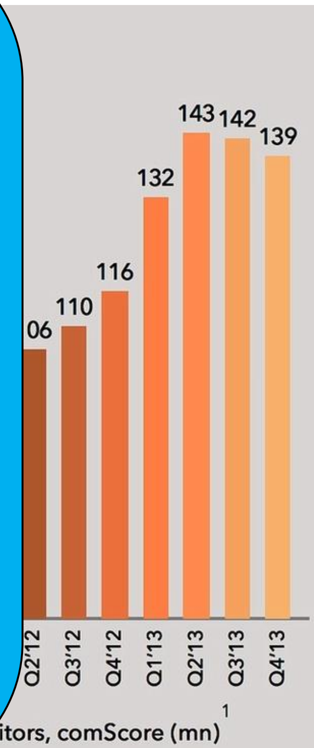




Sources: Vincenzo Cosenza, The State of LinkedIn, http://vincos.it/the-state-of-linkedin/
via Christopher Penn, http://www.shiftcomm.com/2014/02/state-linkedin-social-media-dark-horse/

TUDelft Delft University of Technology

# LinkedIn Is Not Unique: Data Deluge

LinkedIn

400M users

??? edges

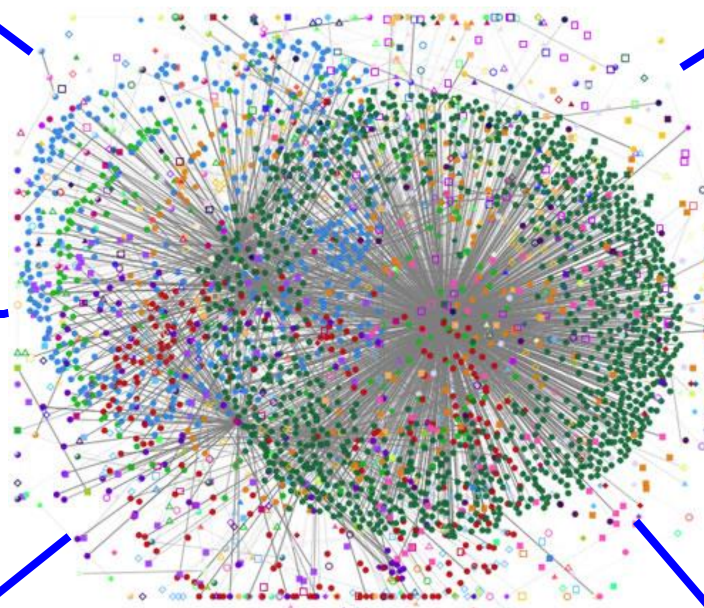

270M MAU
200+ avg followers

>54B edges

1.2B MAU 0.8B DAU
200+ avg followers

>240B edges

# LinkedIn Is Not Unique: Data Deluge

IBM 280k employee-users, 2.6M followers

company/day:
100+ posts, 1,000+ comments

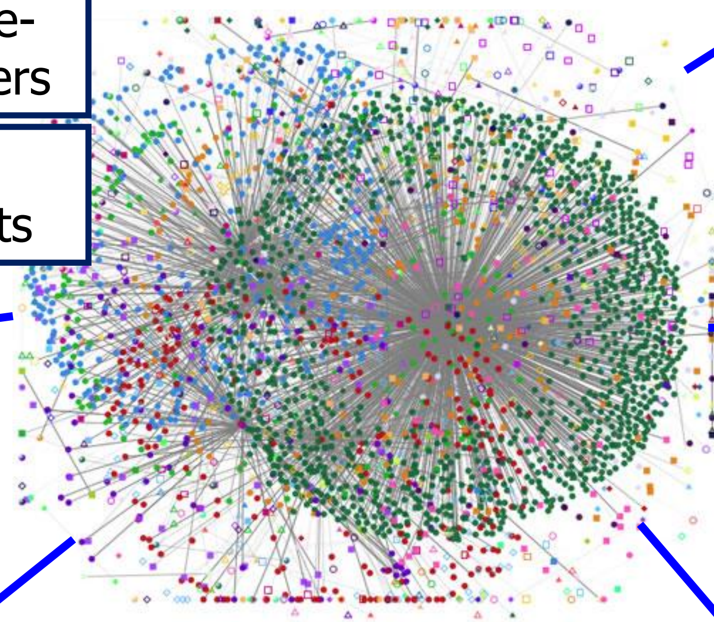

270M MAU
200+ avg followers

>54B edges

1.2B MAU 0.8B DAU
200+ avg followers

>240B edges

# The Data Deluge of Large-Scale Graphs

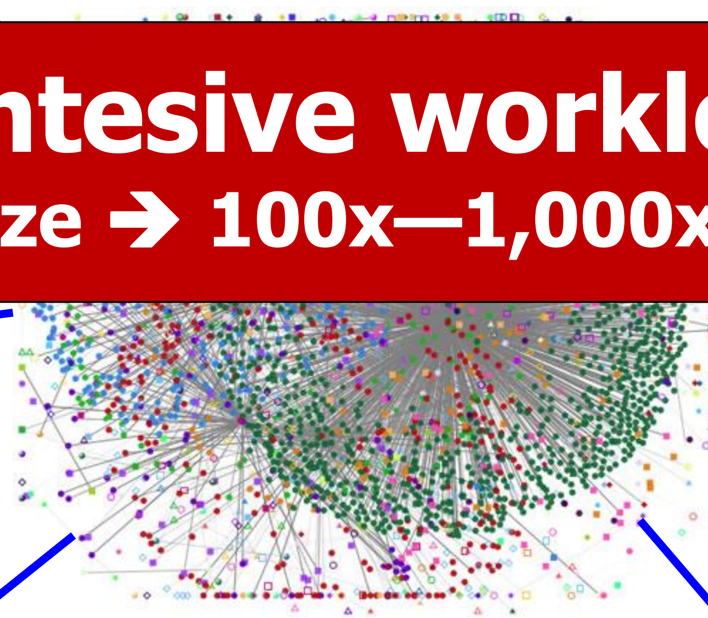

270M MAU

**Data-intesive workload**

**10x graph size → 100x—1,000x slower**

1.2B MAU 0.8B DAU
200+ avg followers

>240B edges

# The Data Deluge of Large-Scale Graphs

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 of Large-Scale Graphs

**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**

TUDelft
Delft University of Technology

LinkedIn Maps

**Ana Lucia Varbanescu's Professional Network**

as of November 28, 2013

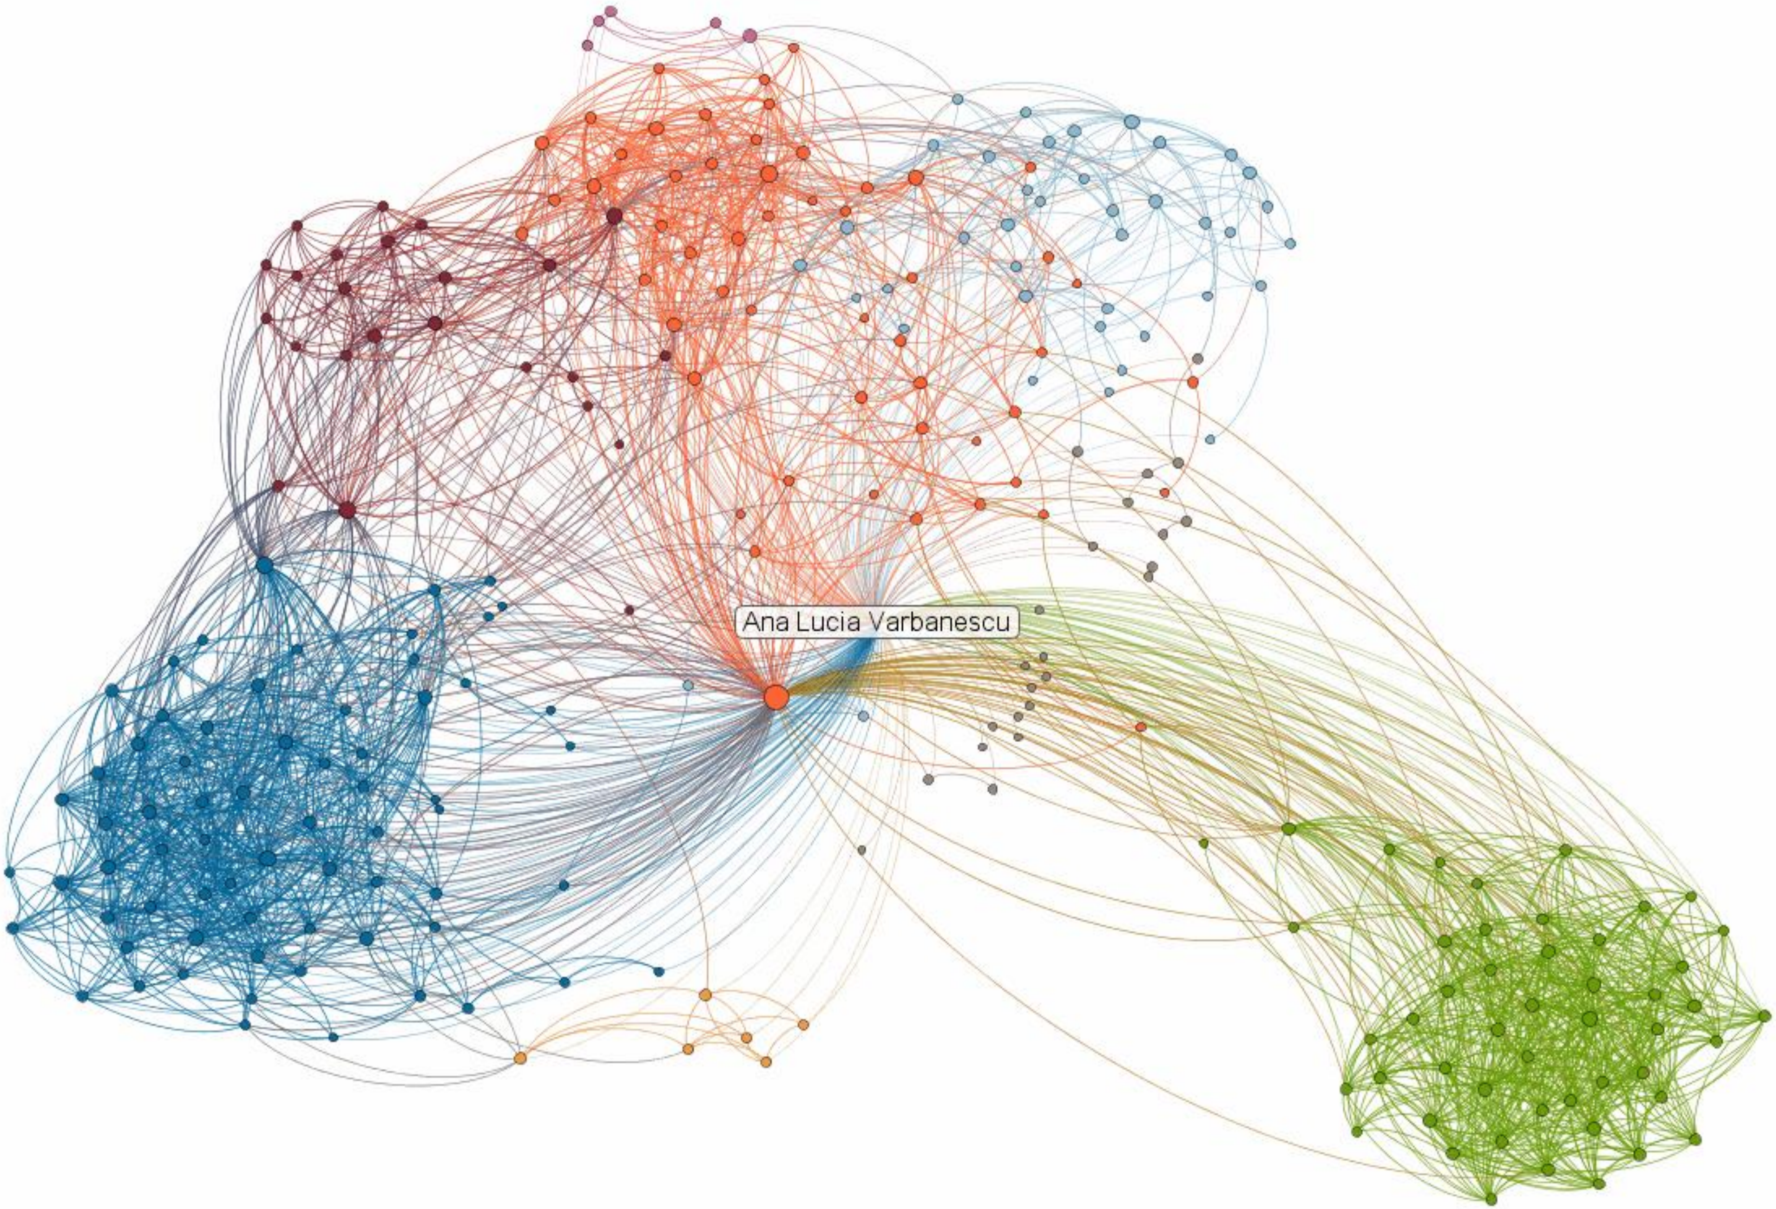

The “sorry, but…” moment

**Supporting multiple users**

**10x number of users → ????x slower**

# Graph Processing @large

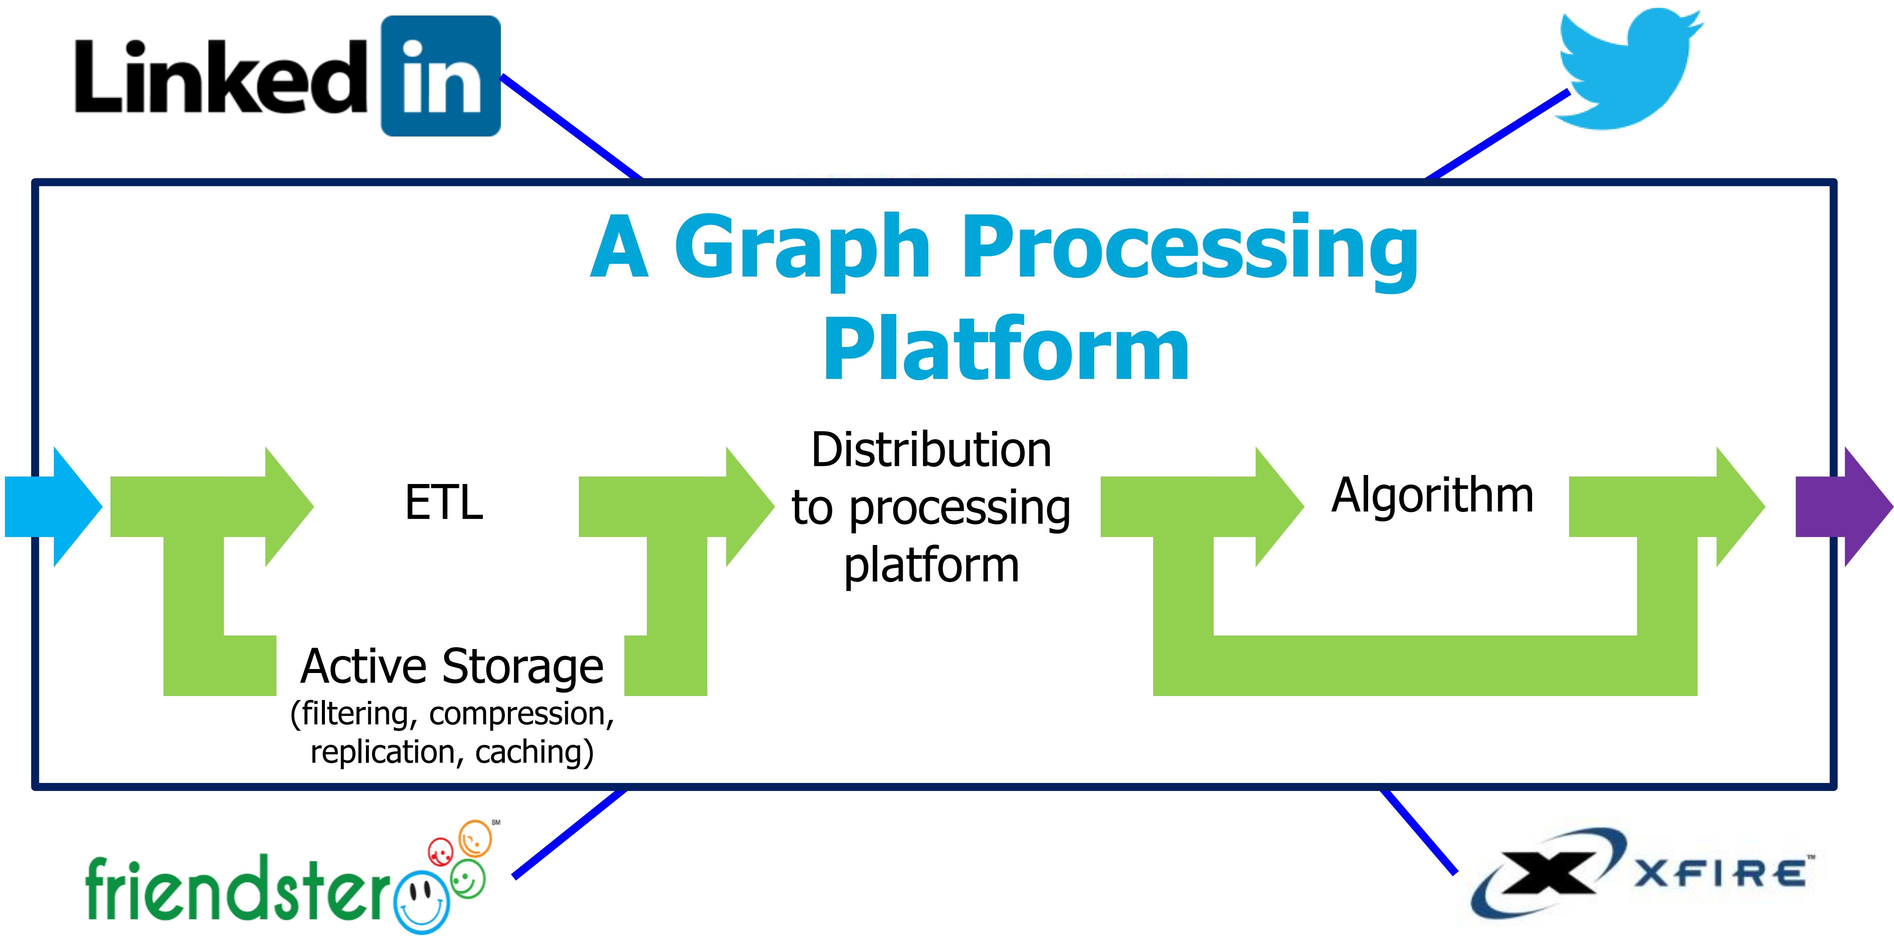


Interactive processing not considered in this presentation.
Streaming not considered in this presentation.

TU Delft
Delft University of Technology

# Graph Processing @large

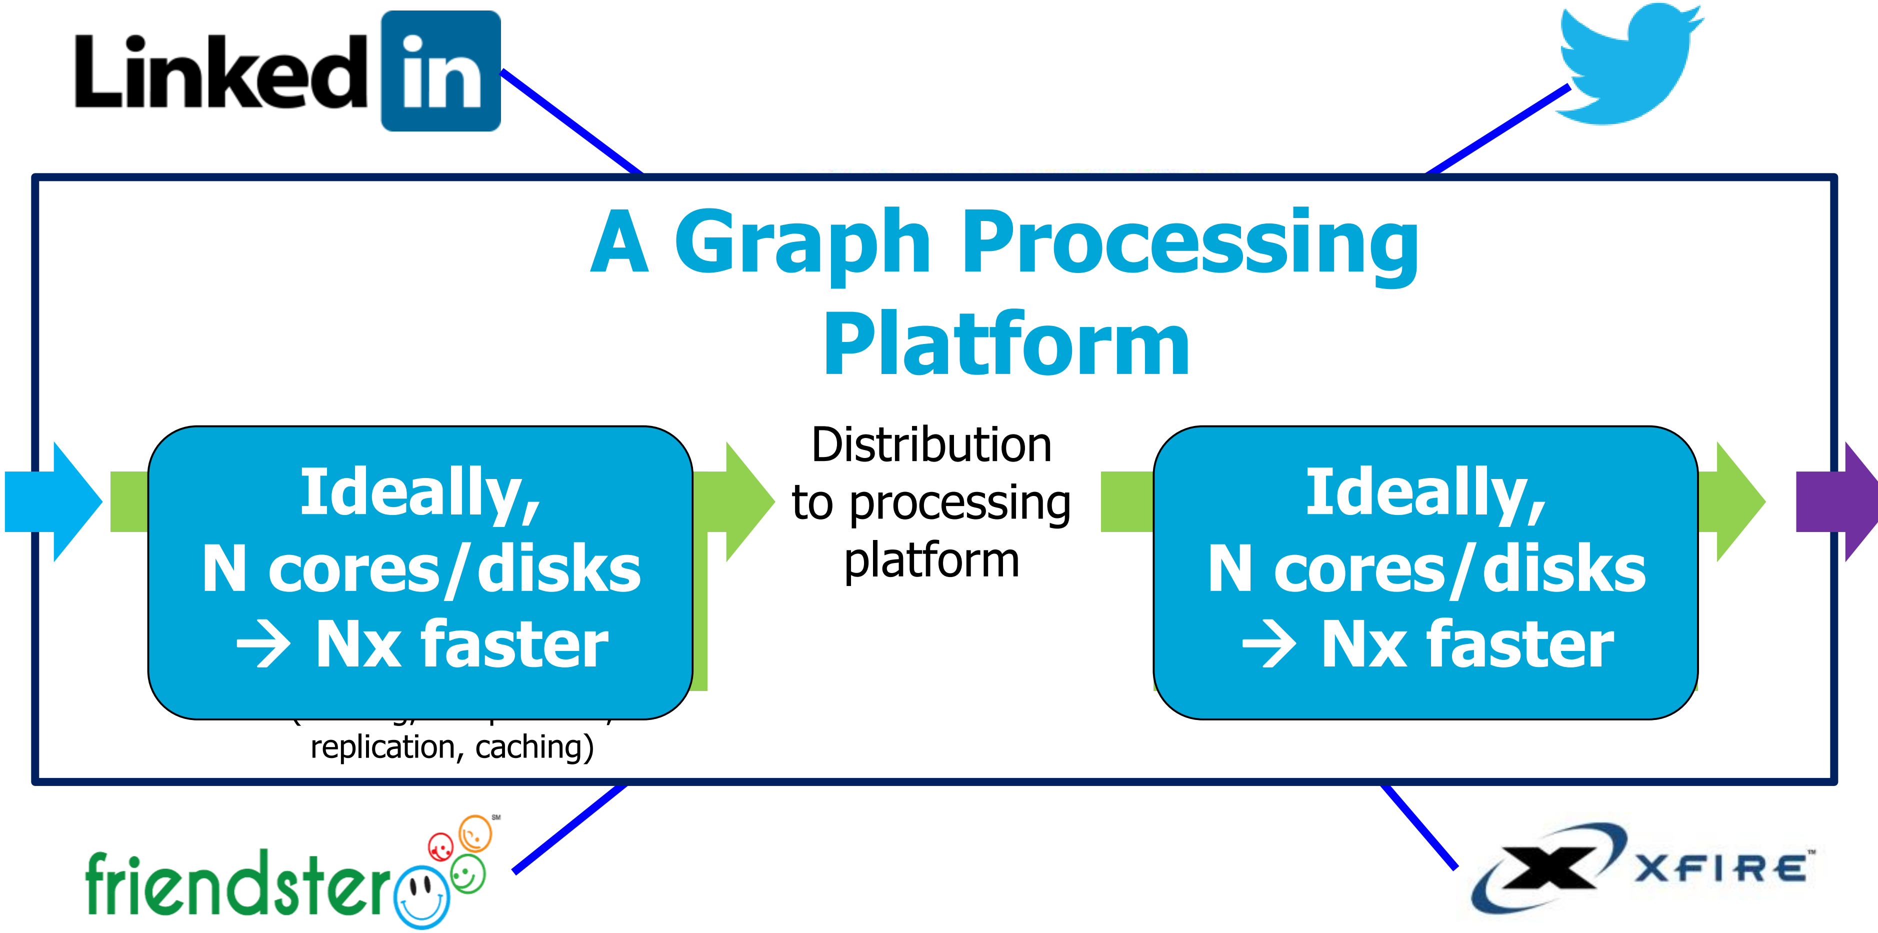


Interactive processing not considered in this presentation.
Streaming not considered in this presentation.

# 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

The graph & RDF benchmark reference

Delft University of Technology

# 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 = Advanced Harness

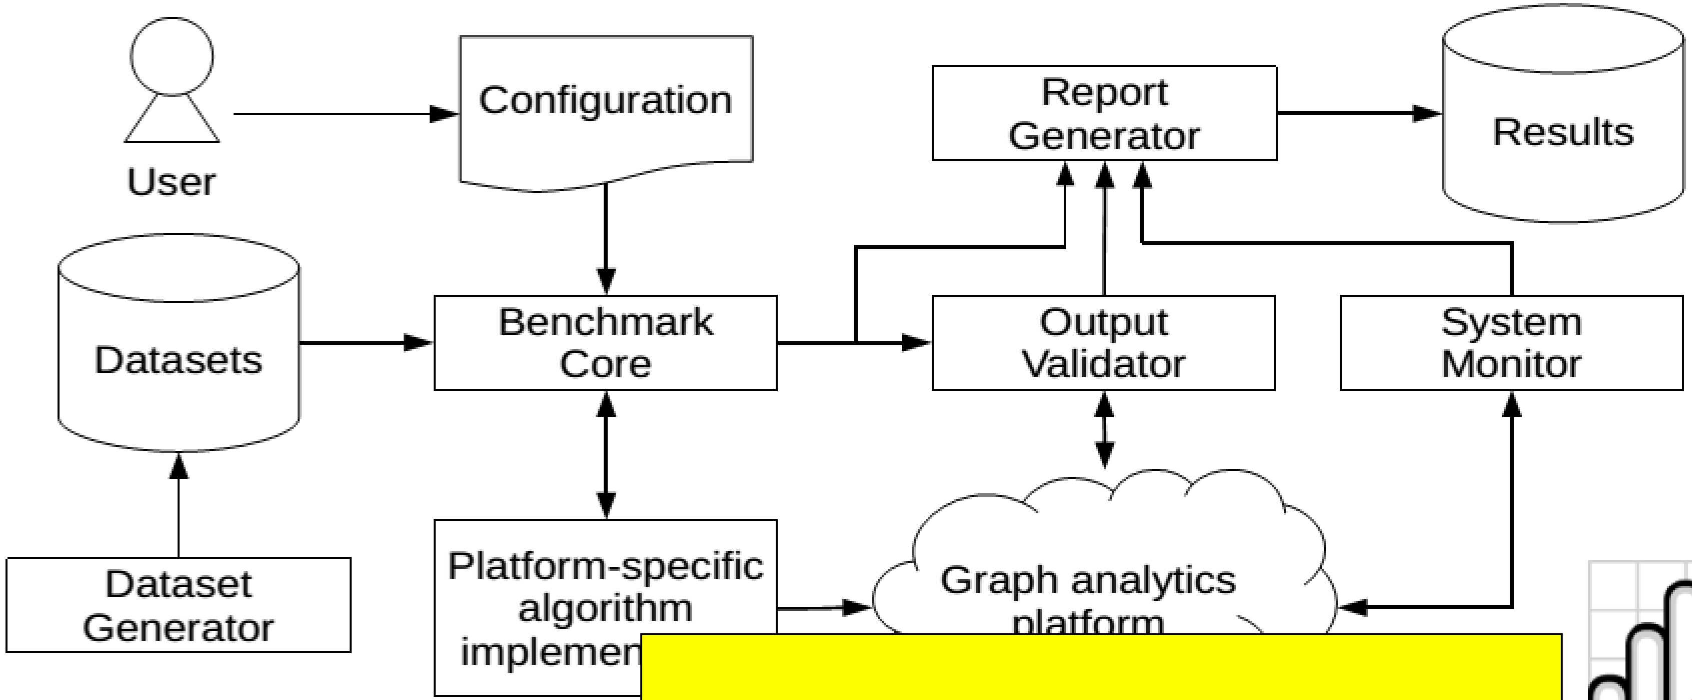


Support of cloud-based platforms technically feasible, methodologically difficult

M. Capota et al., Graphalytics: A Big Data Benchmark for Graph-Processing Platforms. SIGMOD GRADES 2015

# Graphalytics = Real & Synthetic Datasets

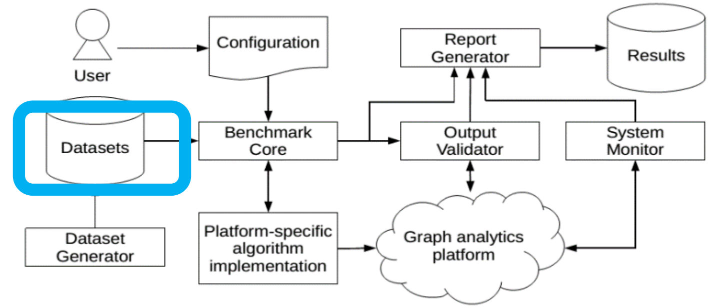


| | Graphs | #V | #E | d | D̄ | Directivity |
|---|---|---|---|---|---|---|
| G1 | Amazon | 262,111 | 1,234,877 | 1.8 | 4.7 | directed |
| G2 | WikiTalk | | | | | directed |
| G3 | KGS | | | | | undirected |
| G4 | Citation | | | | | directed |
| G5 | DotaLeague | 61,171 | 50,870,316 | 2,719.0 | 1,663.2 | undirected |
| G6 | Synth | | | | | undirected |
| G7 | Friendster | | | | | undirected |

Interaction graphs
(possible work)

Property graphs
(planned work)

The Game Trace Archive

https://snap.stanford.edu/ http://www.graph500.org/ http://gta.st.ewi.tudelft.nl/

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

# Graphalytics = Graph Generation w DATAGEN

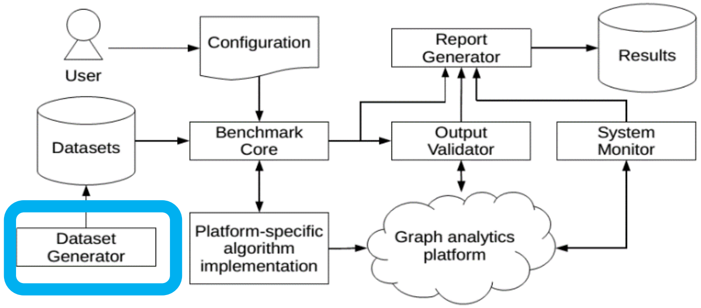


## DATAGEN Process

- Rich set of configurations
- More diverse degree distribution than Graph500
- Realistic clustering coefficient and assortativity

Graphalytics

Person Generation

Edge Generation

"Knows" graph serialization

Activity Generation

Activity serialization

Level of Detail

TU Delft
Delft University of Technology

# Graphalytics = Many Classes of Algorithms

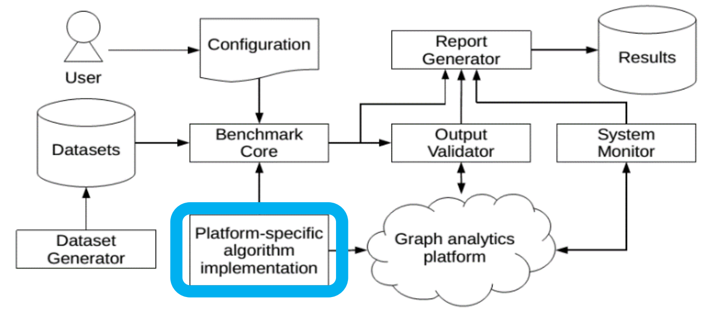


- Literature survey of of metrics, datasets, and algorithms
  - 2009–2013, 124 articles in 10 top conferences: SIGMOD, VLDB, HPDC,

| 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 |
| C | | |

Future work: more diverse algorithms from application domains

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 = Choke-Point Analysis

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

- Examples
  - Network traffic
  - Access locality
  - Skewed execution (stragglers)

**Choke-point analyss often require fine-grained analysis of system operation, across many systems**

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

ongoing work

Delft University of Technology

# Coarse-grained vs Fine-grained Evaluation (1)

## Coarse-grained Method
system viewed as a black-box

Algorithms, Datasets, Resources

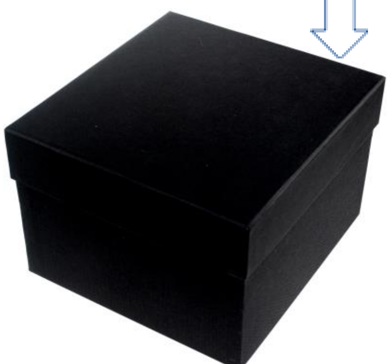

Graph processing system

Coarse-grained performance metrics
(Overall Execution Time)

## Fine-grained Method
system viewed as a white-box

Algorithms, Datasets, Resources

IO operations

Processing operations

Overheads

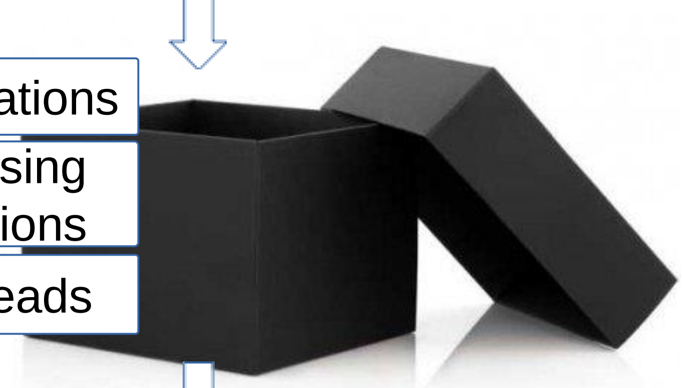

Fine-grained performance metrics
(Stage 3 time, straggler tasks)

Fine-grained evaluation method is more comprehensive

# Coarse-grained vs Fine-grained Evaluation (2)

## Abstract

### Coarse-grained Method
knowledge at conceptual level

Graph Processing Systems

Distributed Infrastructure

several performance results

## Granular

### Fine-grained Method
knowledge at technical level

hadoop

APACHE GIRAPH

GraphLab

Oracle Labs PGX

many performance results

Fine-grained evaluation method is more comprehensive

… but more time-consuming, esp. to implement

# Graphalytics: Granula Overview

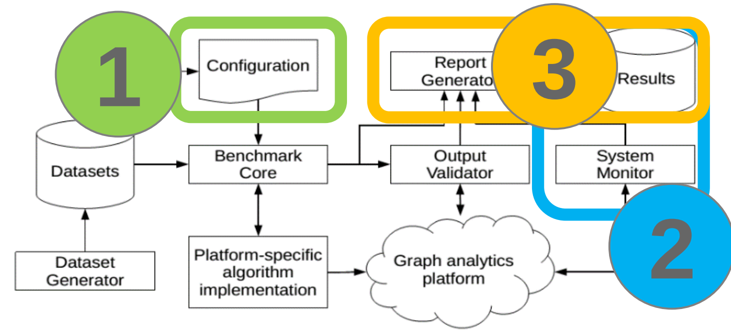


**Granular**

Fine-grained Method

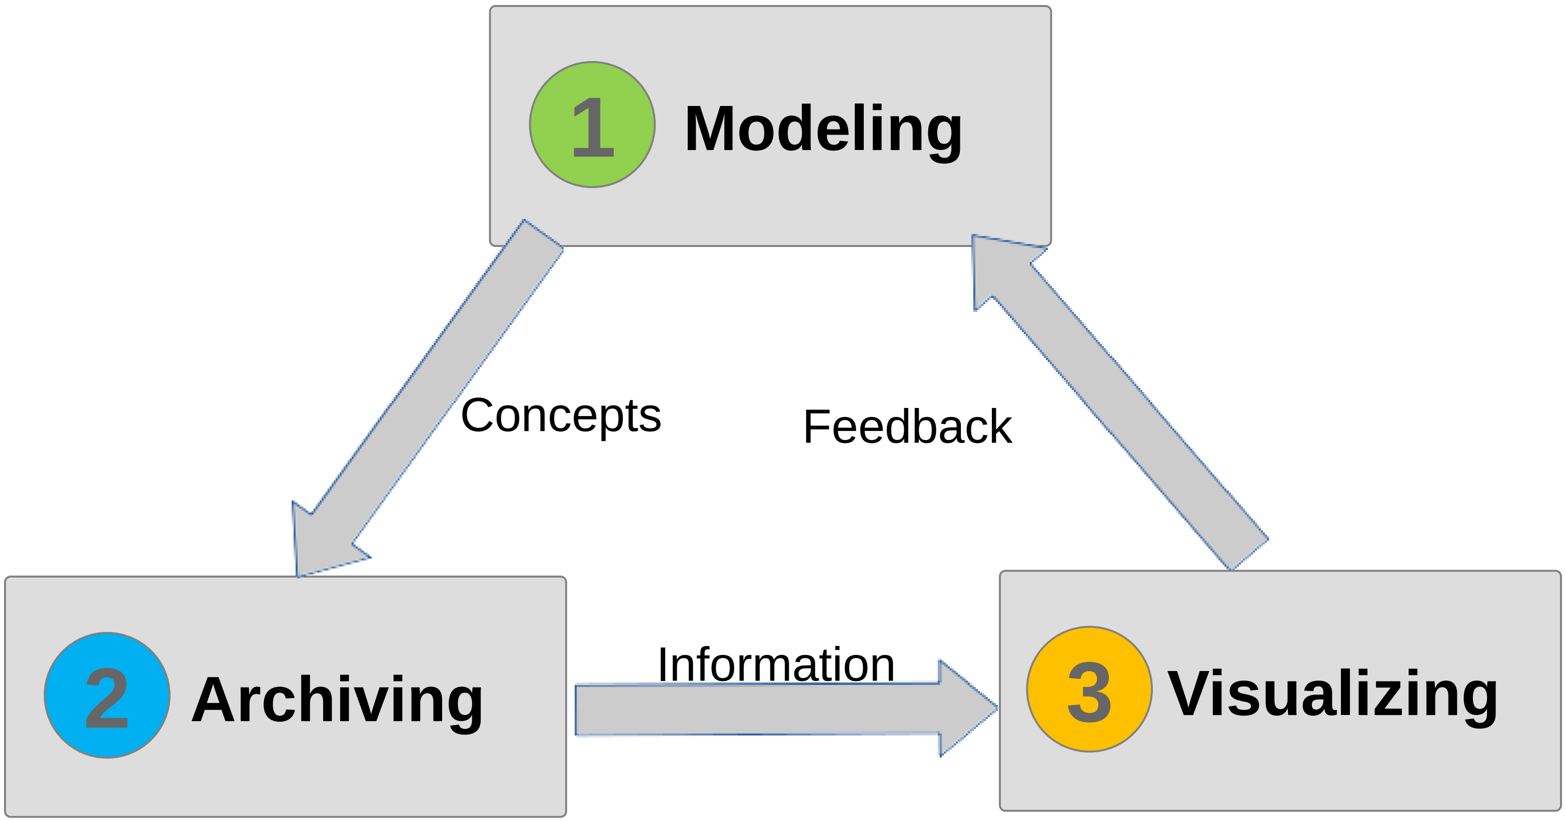


https://github.com/tudelft-atlarge/granula/

TU Delft
Delft University of Technology

# 1 Granula Modeller

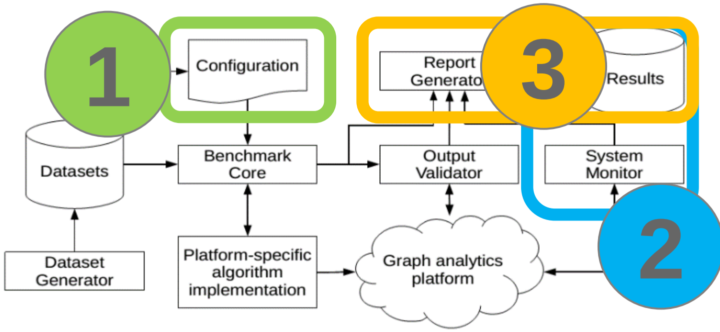


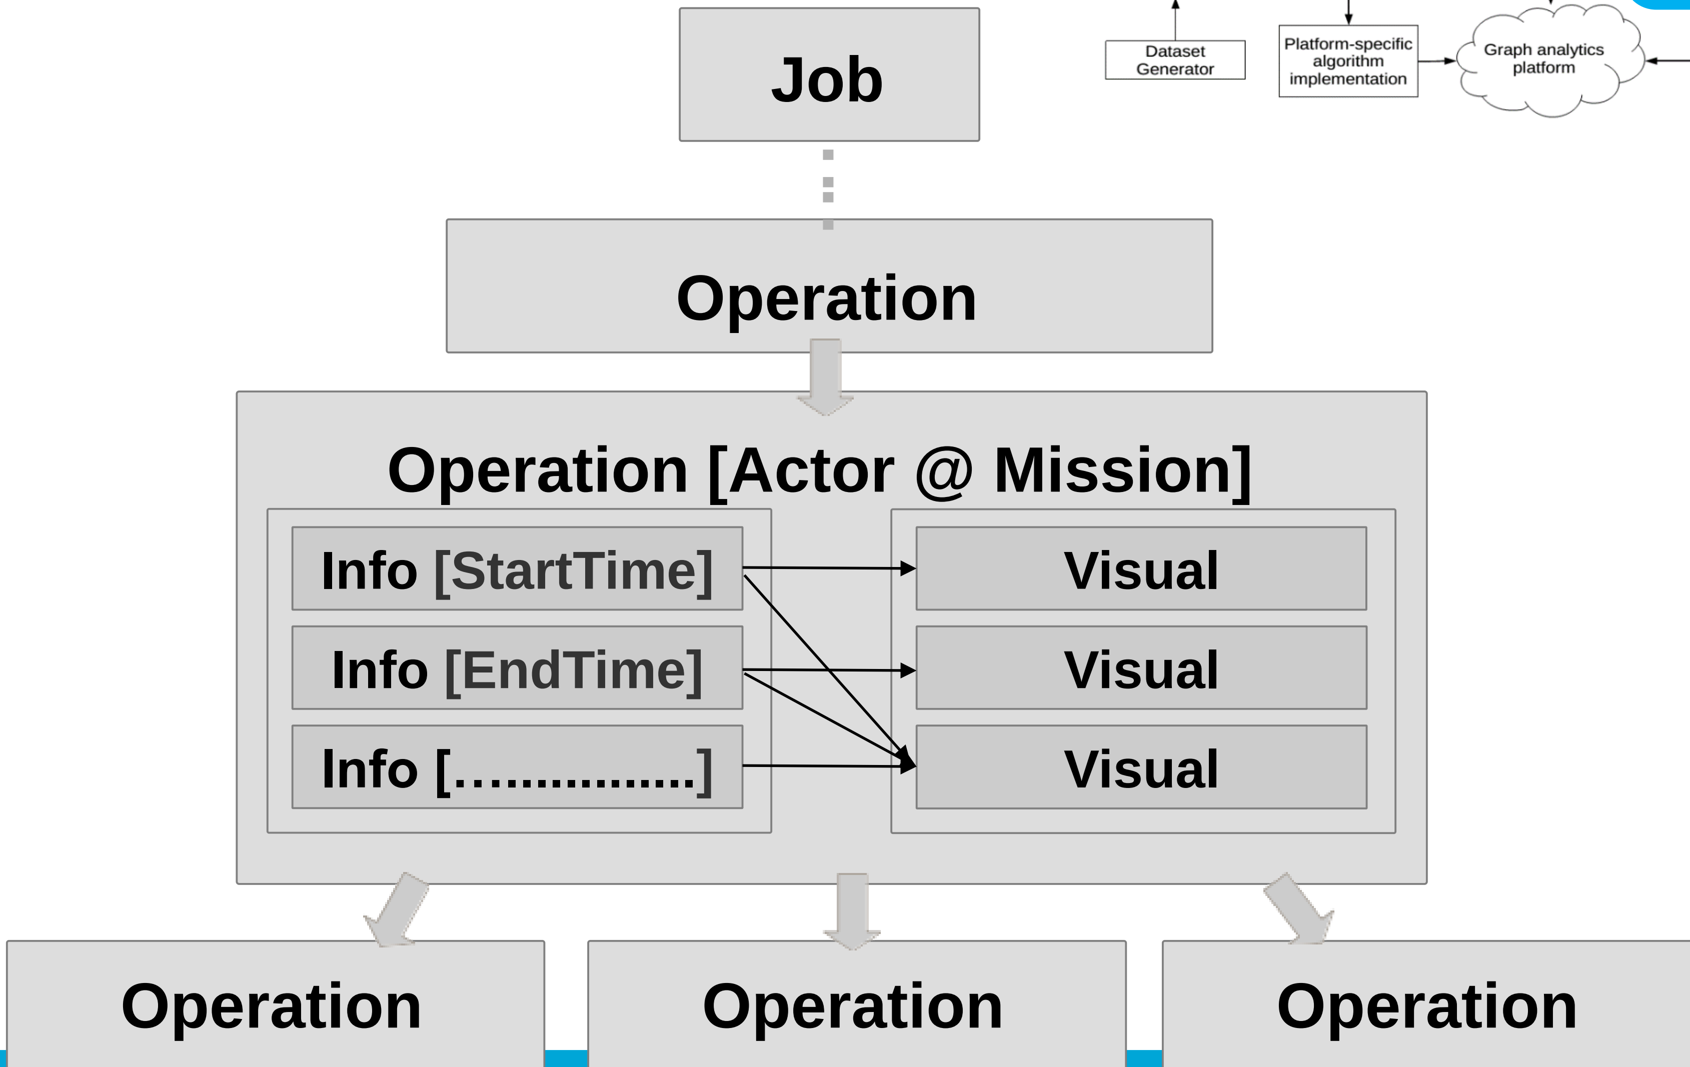


Time-consuming, expert-only, done only once

TU Delft
Delft University of Technology

# 2 Granula Archiver

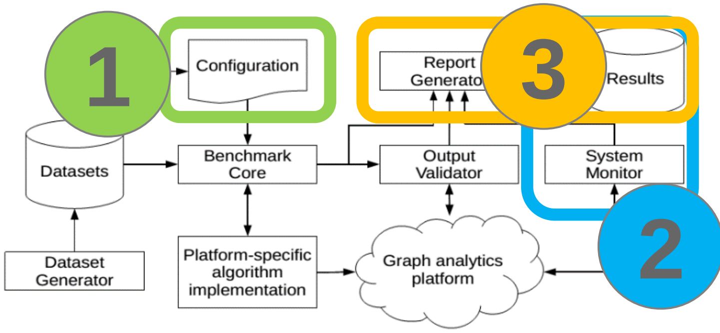


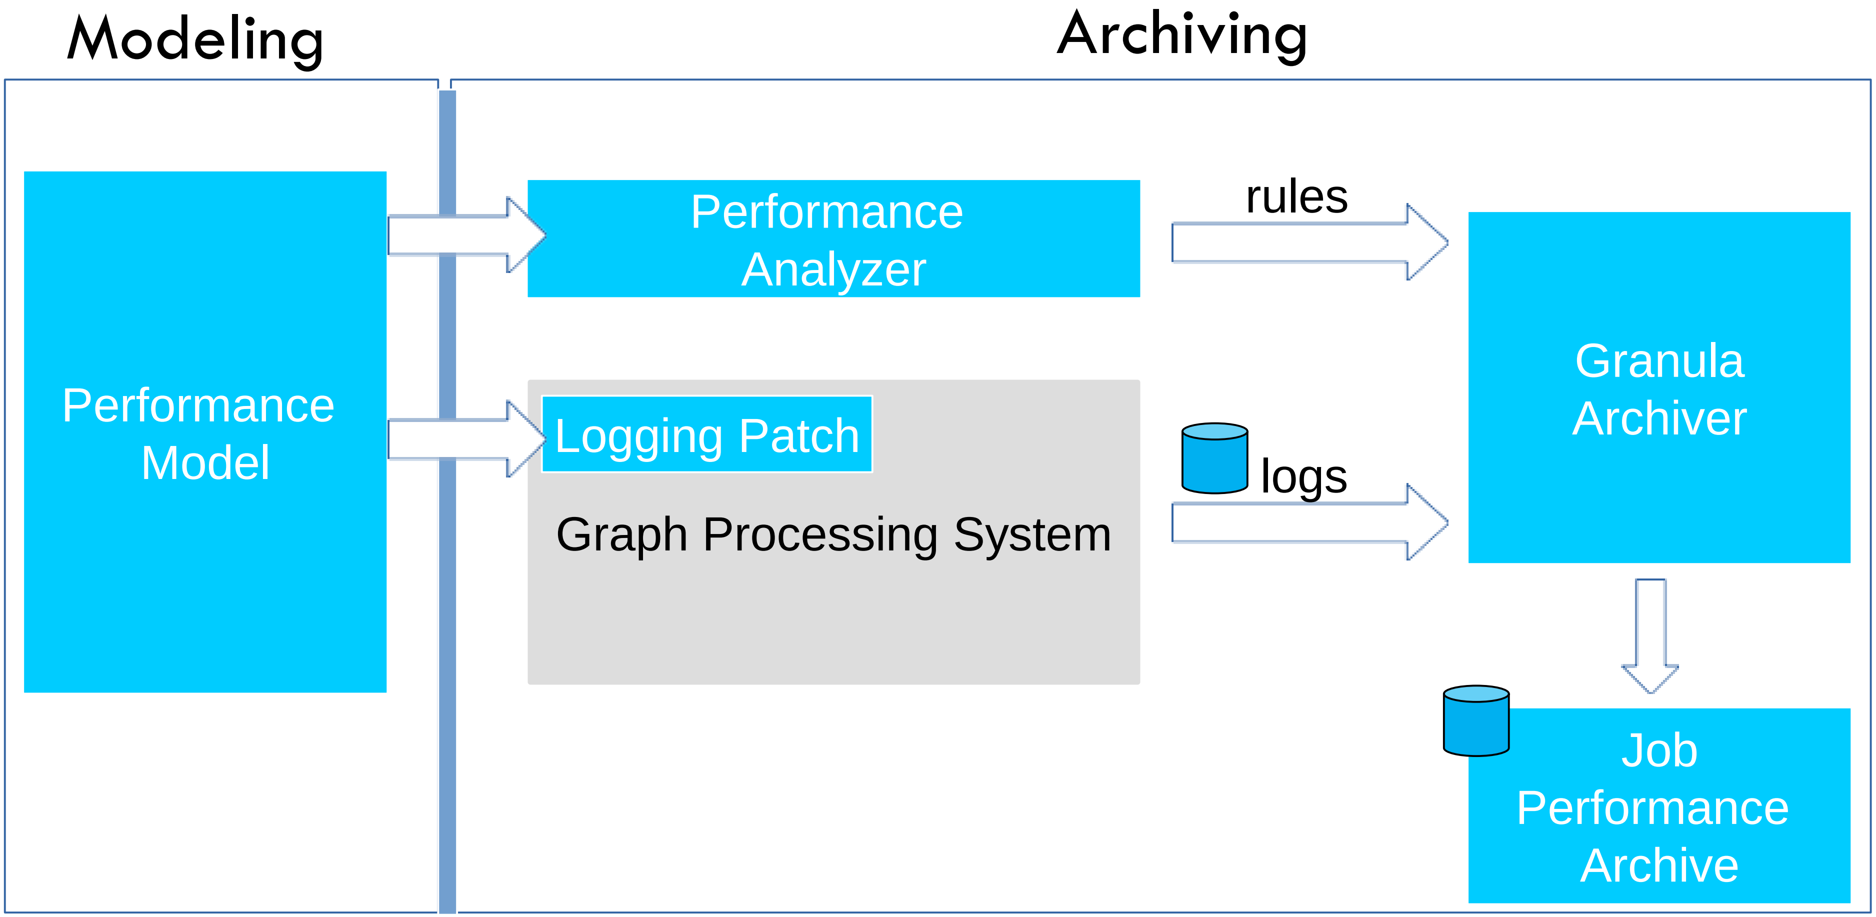


Time-consuming, minimal code invasion,
automated data collection at runtime, portable archive

Delft University of Technology

# 3 Granula Visualizer

Portable choke-point analysis for everyone!

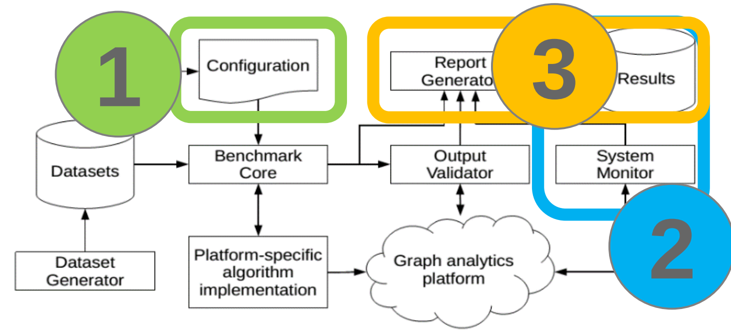


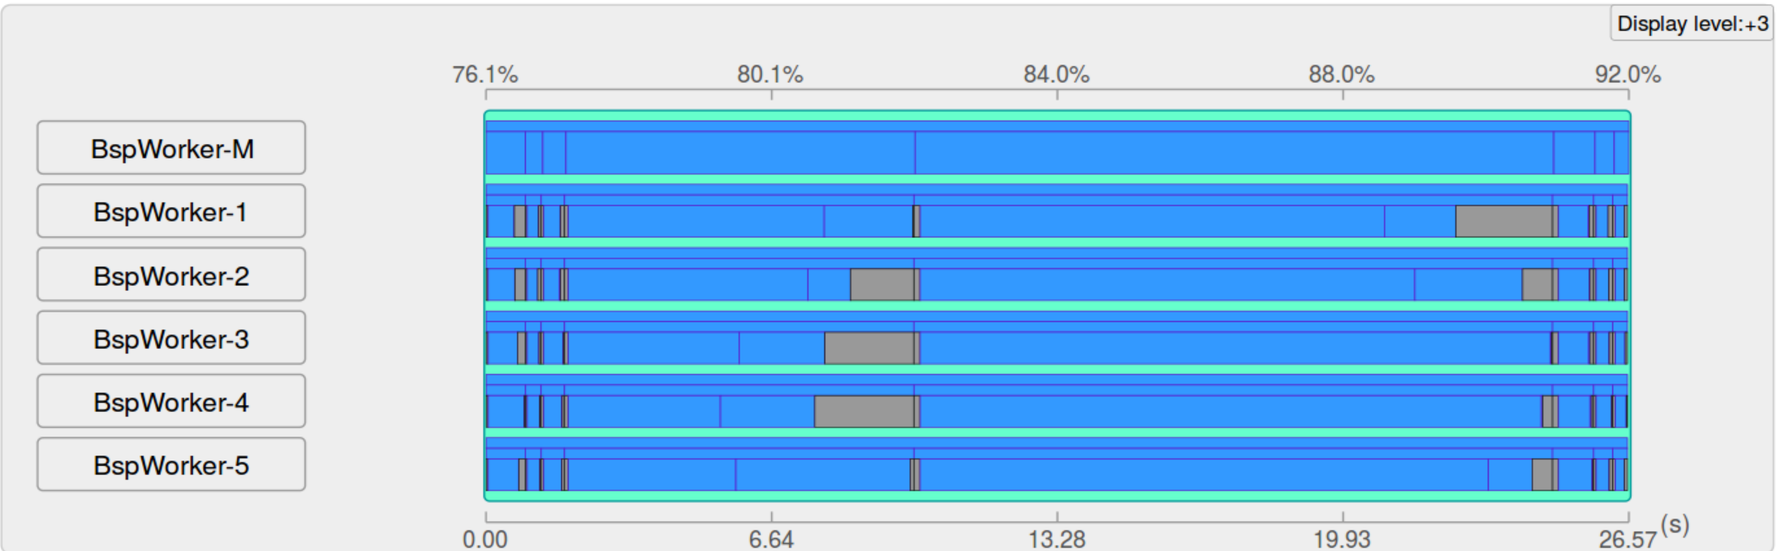


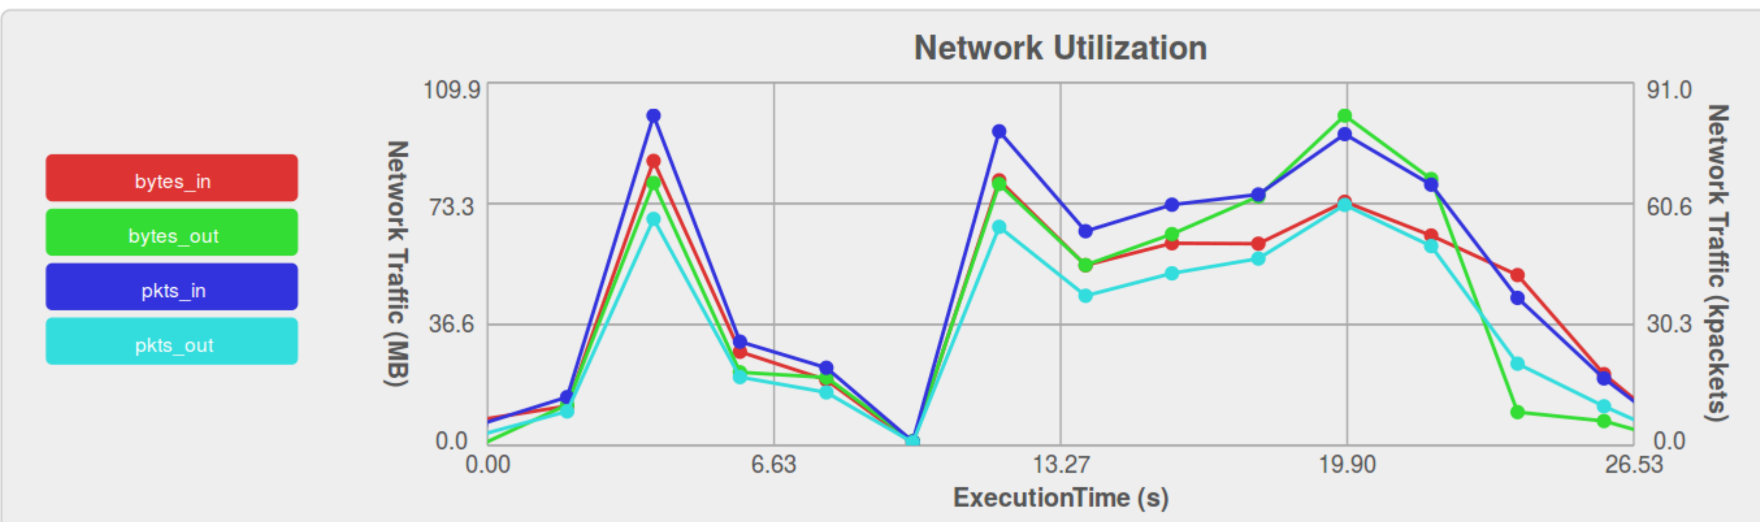

# Graphalytics = Advanced Software Engineering Process

**https://github.com/tudelft-atlarge/graphalytics/**

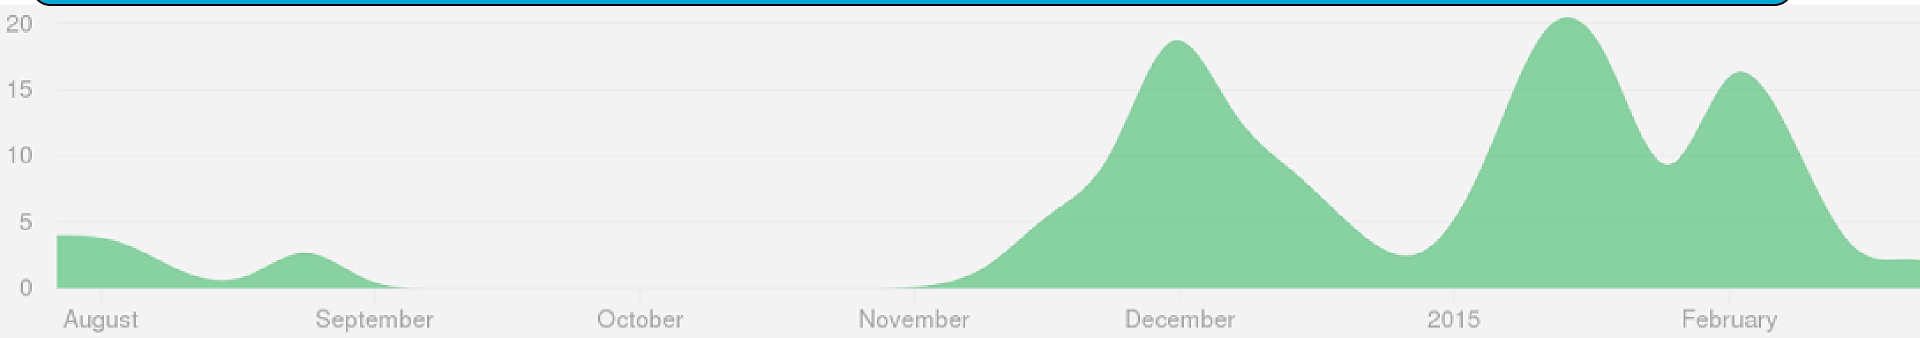


- All significant modifications to Graphalytics are peer-reviewed by developers
  - Internal release to LDBC partners (Feb 2015)
  - Public release, announced first through LDBC (Apr 2015*)
- Jenkins continuous integration server
- SonarQube software quality analyzer

# Graphalytics in Practice

6 real-world datasets +
2 synthetic generators

Data ingestion not included here!

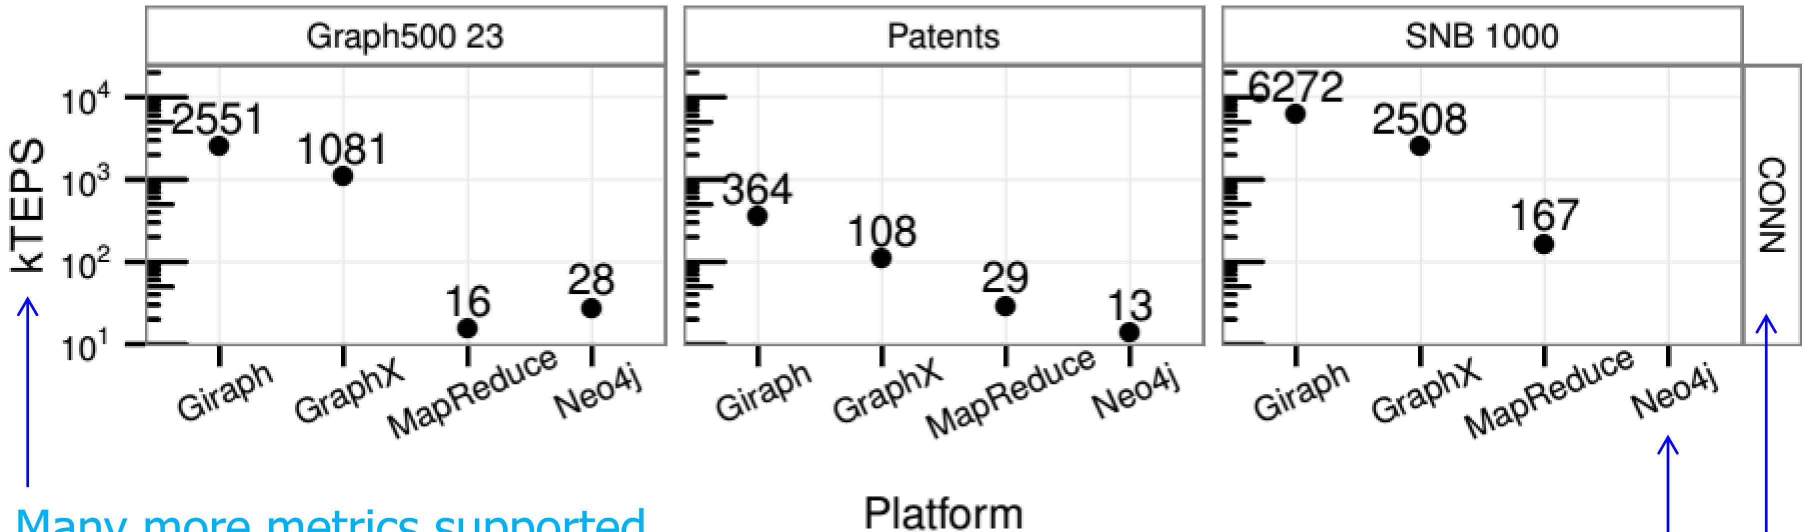


Many more metrics supported

10 platforms tested w prototype implementation

5 classes of algorithms

- Missing results = failures of the respective systems

M. Capota et al., Graphalytics: A Big Data Benchmark for Graph-Processing Platforms. SIGMOD GRADES 2015

TUDelft
Delft University of Technology

# Graphalytics: 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

Guo et al., How Well do Graph-Processing Platforms Perform? IPDPS'14.

Guo et al., An Empirical Performance Evaluation of GPU-Enabled Graph-Processing Systems. CCGRID'15.

# Graphalytics, in a nutshell

- Advanced benchmarking harness
- Diverse real and synthetic datasets
- Many classes of algorithms
- Many classes of algorithms
- Granula for manual choke-point analysis
- Modern software engineering practices
- Supports many platforms
- Ongoing development

LDBC The graph & RDF benchmark reference

DELFT DATA SCIENCE

ORACLE®

bitbrains IT SERVICES

**http://graphalytics.ewi.tudelft.nl**
**https://github.com/tudelft-atlarge/graphalytics/**

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

**http://graphalytics.ewi.tudelft.nl**
**https://github.com/tudelft-atlarge/graphalytics/**

**Join us for the SC2015 tutorial, Nov 15 (tut149)**


Alexandru Iosup
A.Iosup@tudelft.nl



GRAPHALYTICS was made possible by a generous contribution from Oracle.


UPC

CWI

ORACLE

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

LDBC The graph & RDF benchmark reference

\+ many other contributors

DELFT DATA SCIENCE

TUDelft Delft University of Technology

# A few extra slides

# Discussion

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

**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)

Delft University of Technology

# Graphs at the Core of Our Society: The LinkedIn Example→Data Deluge




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→Data Deluge

## The State of **LinkedIn**




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/