Distributed Systems Group

Department Software Technology

Prof.dr.ir. Dick H.J. Epema

September 2019

TU Delft Delft University of Technology

TU Delft

Delft University of Technology

The Distributed Systems Group People

Prof. Dick Epema

scheduling and resource management blockchain

Dr. Lydia Chen

robust, slim and private machine learning systems

Dr. Johan Pouwelse

blockchain trust in the internet

Dr. Jan Rellermeyer

middleware for big-data processing

Dr. Stefanie Roos

scalability of blockchains anonymous communication networks

TU Delft

The Distributed Systems Group Teaching

MSc courses:

- CS4215: Quantitative Performance Analysis of Computer Systems (Q1)

- IN4150: Distributed Algorithms (Q2)

  (core Software Technology)

- IN4391: Distributed Computing Systems (Q3)

  (core Data Science & Technology)

- IN4253: Blockchain Engineering (Q3)

- IN4392: Seminar Cloud Computing (Q4)

TU Delft

The Distributed Systems Group
Research Flavor

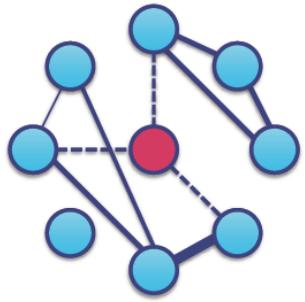

- Our research is:
  - **fundamental**: devise new application-independent concepts in distributed systems
  - **experimental**: show the value of new concepts in prototypes or real deployments

TU Delft

The Distributed Systems Group Research Topics

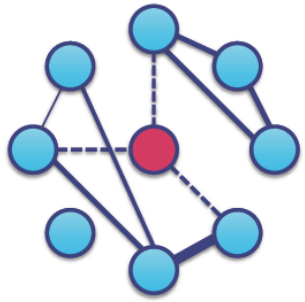

1. Resource Management and Scheduling
2. Big Data Processing
3. Cooperative Systems (trust, blockchain)

See research pages and annual report 2015-2018 at

http://www.ds.ewi.tudelft.nl

Experimentation: DAS6 on the way!!

VU (136 CPUs)

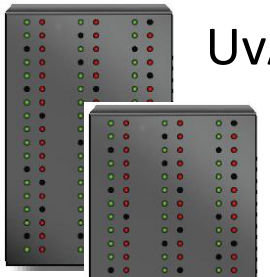

UvA/MultimediaN (62)

UvA (36)

SURFnet6

- **System purely for CS research**
- **Operational since June 2015**
- **Specs:**
  - 3,200 cores (8-core CPUs)
  - 2.4 GHz CPUs
  - accelerators (GPUs)
  - 800 TB storage
  - 10 Gb/s Infiniband
  - Gb Ethernet

10 Gb/s lambdas

TU Delft (48)

Astron (9)

Leiden (48)

- **Article in IEEE Computer 49(5), 2016**

Advanced School for Computing and Imaging

Netherlands Organisation for Scientific Research

TU Delft

The KOALA multicluster scheduler

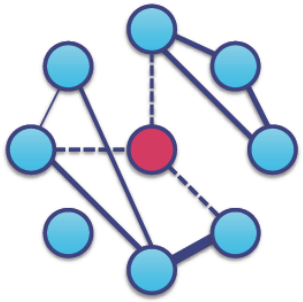

deployed
on the DAS
since 2005

parallel applications
workflows
MapReduce / Spark

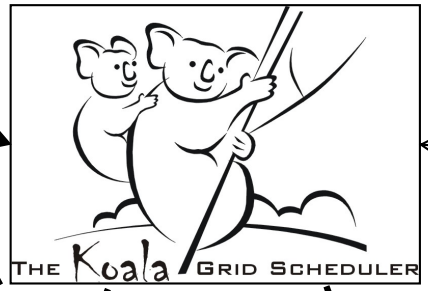

information service

Local
Resource
Managers

LRM

LRM

LRM

LRM

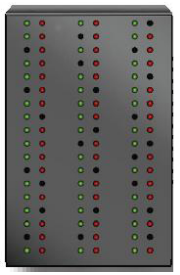

TU Delft

Big Data Processing (1/2)

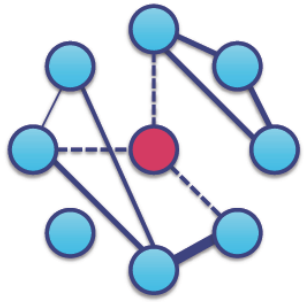

graph processing

- benchmarking
- performance analysis

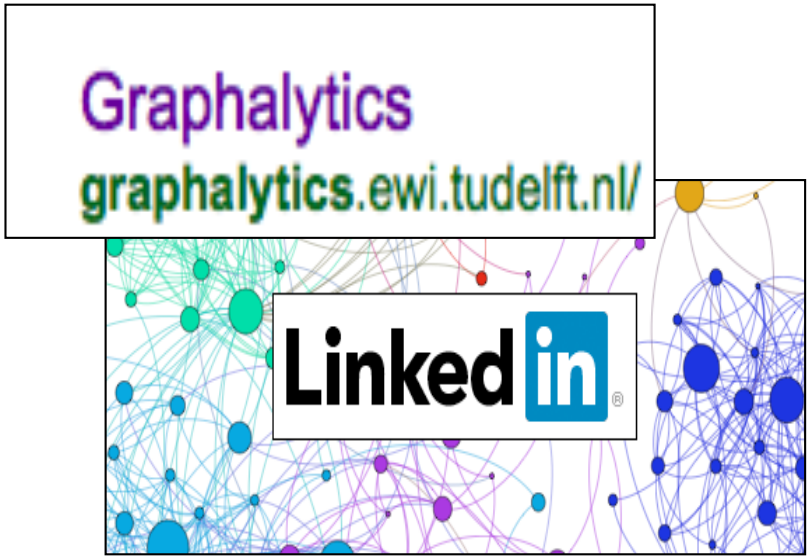

data processing frameworks:

- optimizations for new hardware
- forget about VMs

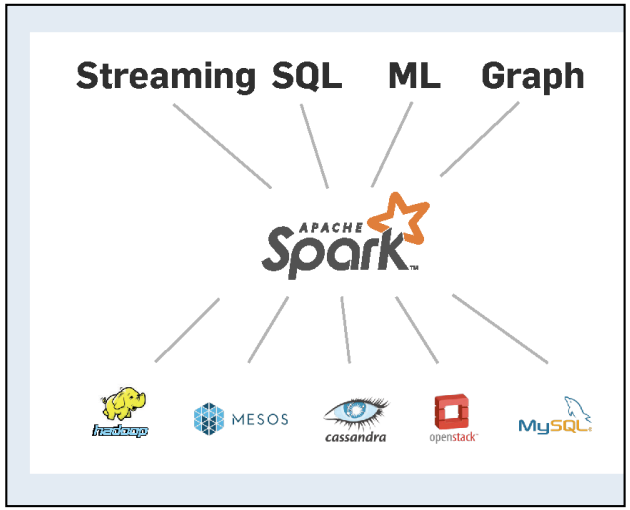

Jan Rellermeyer

TU Delft

Big Data Processing (2/2): efficient, robust, private, fair learning systems

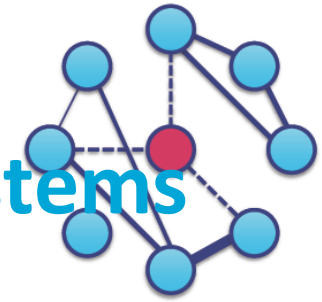

Processing Systems
--anomaly detection
--sprinting
--tail latency
--dependability
--workload analysis

Learning System
--slim and private
--robust and adversarial
--large scale and efficiency
--novel applications

Artificial Intelligence
--active learning
--fair learning
--distributed learning

Lydia Chen

Cooperative Systems (1/3): Tribler

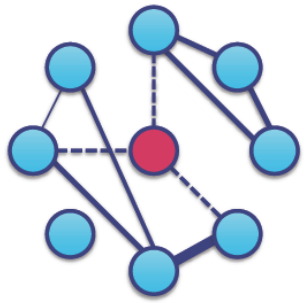

- Is based on the **BitTorrent P2P** file-sharing system
- Uses an **epidemic protocol** for peer and content discovery
- Was **first released** on 17 March 2006 (2,000,000+ downloads)
- Enables **video-on-demand** and **live streaming**
- Is our **research vehicle** for P2P, Online Social Networks, reputation systems, blockchain

tribler

- **Current focus**: reputation, trust, blockchain
- **Download** at www.tribler.org

Johan Pouwelse

Cooperative Systems (2/3): trust/reputation

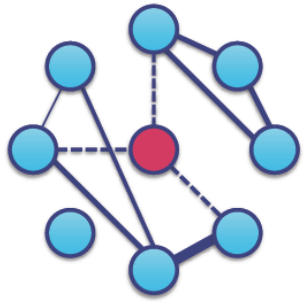

- **Problems:**
  - why help others downloading in P2P systems?
  - why contribute to Wikipedia?
  - why trust money without central banks?

- **Solution:** create a trust system **without central control**
  - record decentralized interaction history
  - disseminate this history in an attack-resilient way
  - create decentralized markets

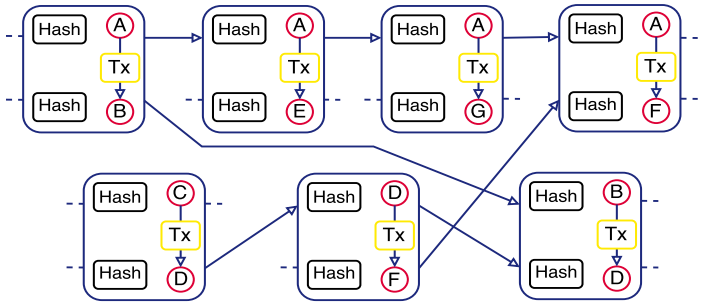

- **In Tribler:**
  - **Trustchain:** alternative to the blockchain

TU Delft

Cooperative Systems (3/3): Anonymity and Blockchain

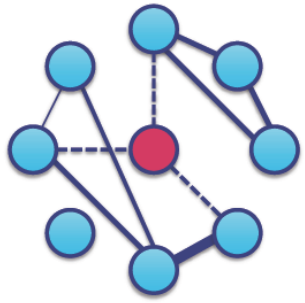

Scalable Anonymity

- How to deal with millions of Tor users?
- More efficient protocols, incentives to contribute, …

Anonymity and Blockchain

- Attacking Zcash's and Monero's anonymity
- Building network layer protocols for blockchain anonymity

Stefanie Roos

More information

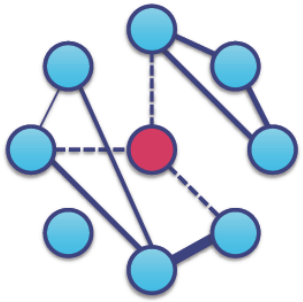

- **MSc coordinator:** Jan Rellermeyer
- **Some previous MSc theses:**
  - www.ds.ewi.tudelft.nl/epema/teaching
- **Home page Distributed Systems**:
  - www.ds.ewi.tudelft.nl
- **Web sites projects:**
  - KOALA: www.st.ewi.tudelft.nl/koala
  - DAS5: www.cs.vu.nl/das5
  - Tribler: www.tribler.org

TU Delft

Distributed Systems Tag Cloud

TU Delft