MASSIVIZING COMPUTER SYSTEMS

Co-sponsored by:

NWO Netherlands Organisation for Scientific Research

STW connecting innovators

= MAKING COMPUTER SYSTEMS SCALABLE, RELIABLE, PERFORMANT, ETC., YET ABLE TO FORM AN EFFICIENT ECOSYSTEM

@Large Research
Massivizing Computer Systems

http://atlarge.science

Many thanks, Arie, Andy, and In-Vivo Analytics for Big Software Quality in general.
Thanks also to the Lorentz Center team in Leiden, the Netherlands.

VU VRIJE UNIVERSITEIT AMSTERDAM

bit.ly/AI18LorentzTalk

@Alosup

Prof. dr. ir. Alexandru Iosup

MASSIVIZING COMPUTER SYSTEMS

Co-sponsored by:

NWO Netherlands Organisation for Scientific Research

STW connecting innovators

= MAKING COMPUTER SYSTEMS SCALABLE, RELIABLE, PERFORMANT, ETC.; YET ABLE TO FORM AN EFFICIENT ECOSYSTEM

DISTRIB.SYS. X PERF.ENG. X SW.ENG.

@Large Research
Massivizing Computer Systems

http://atlarge.science

Many thanks, Arie, Andy, and In-Vivo Analytics for Big Software Quality in general.
Thanks also to the Lorentz Center team in Leiden, the Netherlands.

VU VRIJE UNIVERSITEIT AMSTERDAM

bit.ly/AI18LorentzTalk

@Alosup

Prof. dr. ir. Alexandru Iosup

ATLARGE RESEARCH, OUR TEAM

http://atlarge.science/people.html

- Professor
- Assistant Prof.
- Teacher
- Post-doc
- Ph.D. student
- Scientist

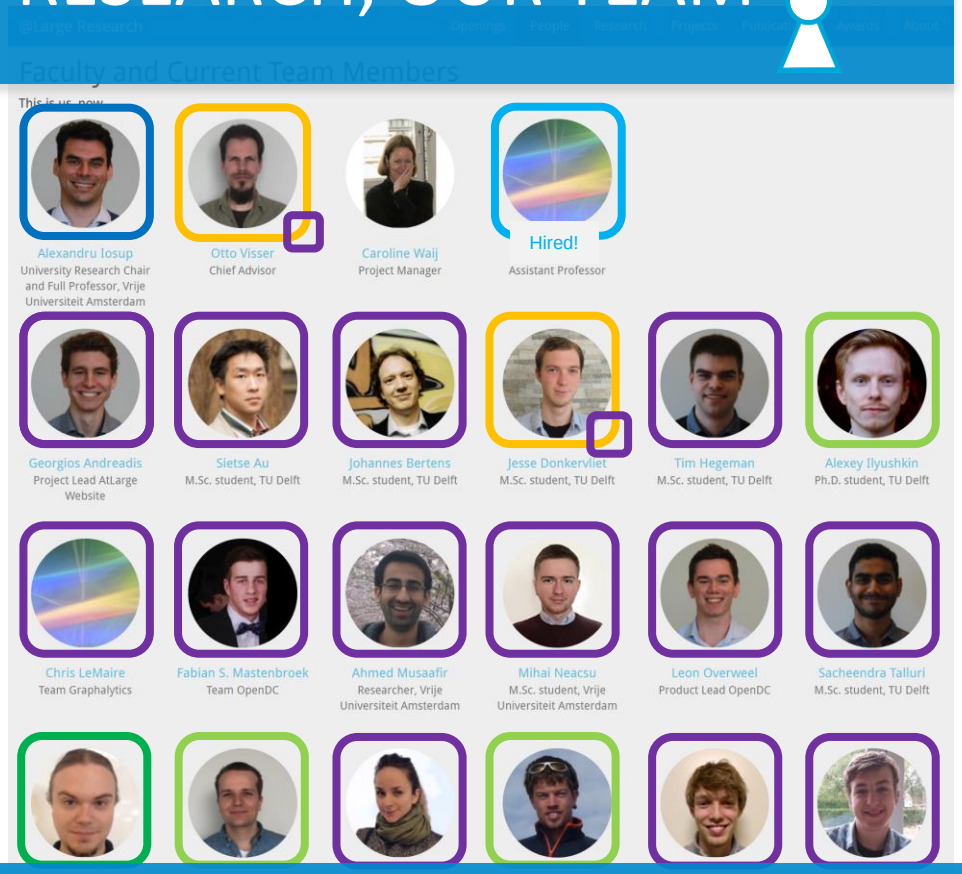

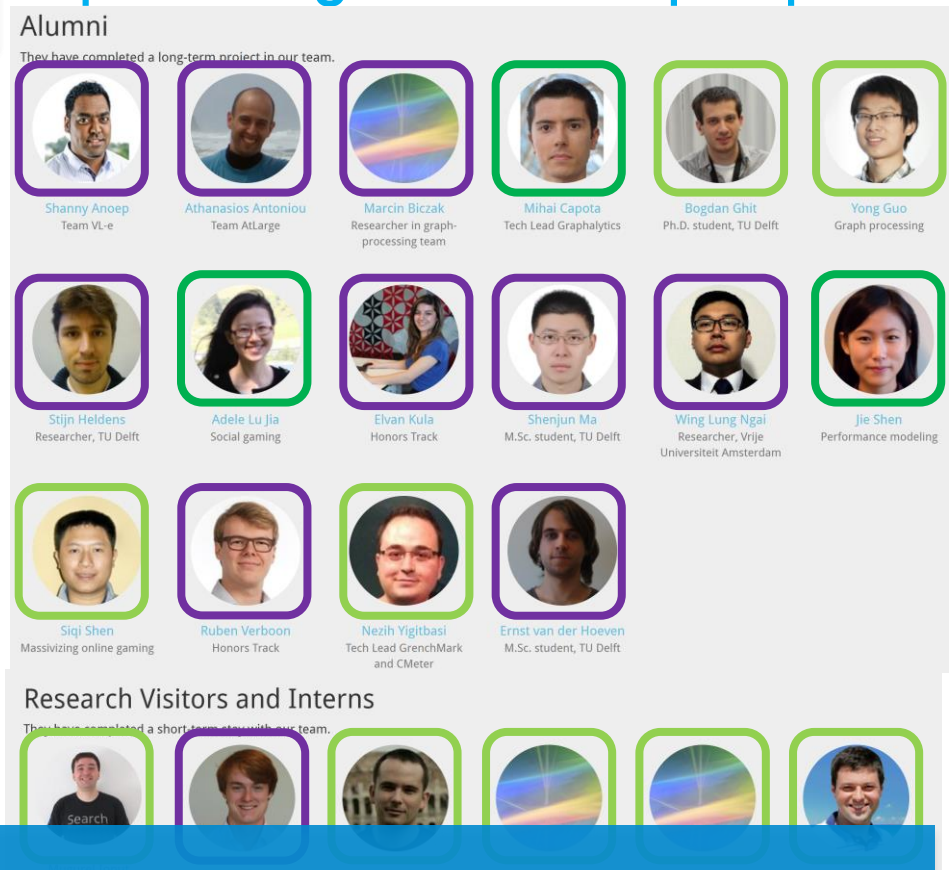

WE ARE A FRIENDLY, DIVERSE GROUP, OF DIFFERENT RACES AND ETHNICITIES, GENDERS AND SEXUAL PREFERENCES, VIEWS OF CULTURE , POLITICS, AND RELIGION. YOU ARE WELCOME TO JOIN!

WHO AM I?
PROF. DR. IR. ALEXANDRU IOSUP

- Education, my courses:
  - > Systems Architecture (BSc)
  - > Distributed Systems, Cloud Computing (MSc)
- Research, 15 years in DistribSys:
  - > Massivizing Computer Systems

- About me:
  - > Worked in 7 countries, NL since 2004
  - > I like to help… I train people in need
  - > VU University Research Chair + Group Chair
  - > NL ICT Researcher of the Year
  - > NL Higher-Education Teacher of the Year
  - > NL Royal Young Academy of Arts & Sciences

VU

MASSIVIZING COMPUTER SYSTEMS: OUR MISSION

1. Improve the lives of millions through impactful research.

2. Educate the new generation of top-quality, socially responsible professionals.

3. Make innovation available to society and industry.

VU VRIJE UNIVERSITEIT AMSTERDAM

http://atlarge.science/about.html

Solvinity. This is IT
Nerdalize
spec Research
LDBC
salesforce
databricks
intel
netherlands eScience center
Amsterdam Data Science
DELFT DATA SCIENCE
ORACLE
IBM
SAP

~60’

Massivizing Computer Systems
A Structured Discussion

~3’ — About Our Team

~30’ — The Golden Age of Distributed Ecosystems … and a Crisis

Interrupts welcome

The main challenges

How we address them: Massivizing Computer Systems

~25’ — Massivizing Computer Systems: Let’s Collaborate

Key for our discussion

What can DistribSys x PerfEng x SwEng do together?

8 ideas for collaboration

~2’ — Take-Home Message

With further reading

THIS IS THE GOLDEN AGE OF DISTRIBUTED ECOSYSTEMS

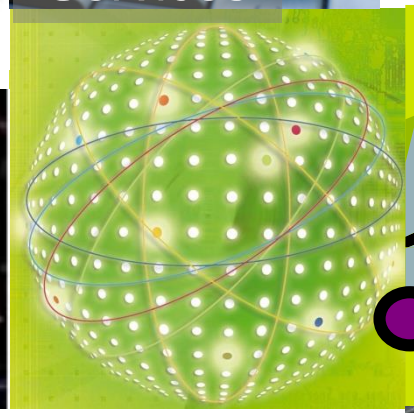

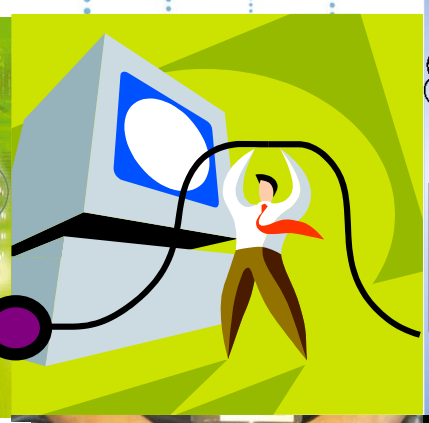

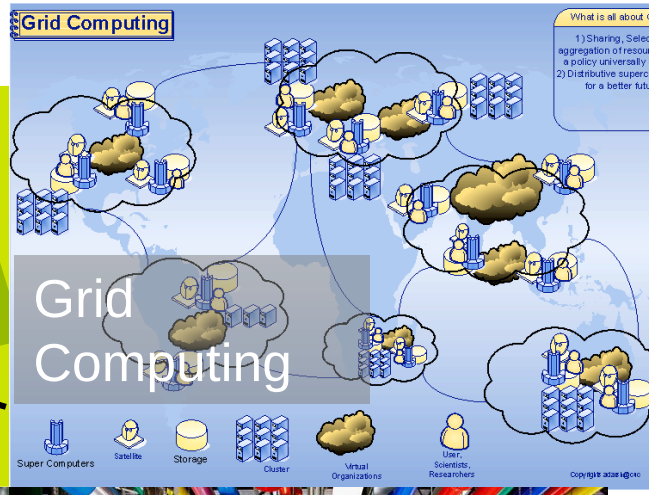

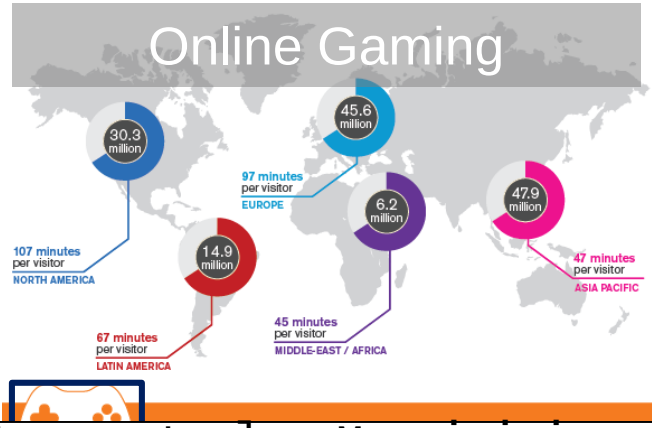

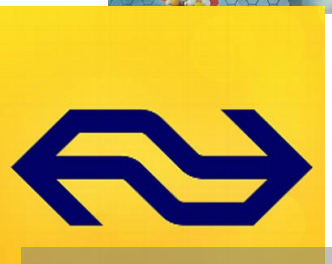

Iosup et al., Massivizing Computer Systems, ICDCS 2018. [Online]

THIS IS THE GOLDEN AGE OF DISTRIBUTED COMPUTER SYSTEMS

Do you recognize this App?

Here is how it operates…

Daily Life

THE CURRENT TECHNOLOGY STACK: DATACENTER, SCHEDULER

Creators

Digital Services

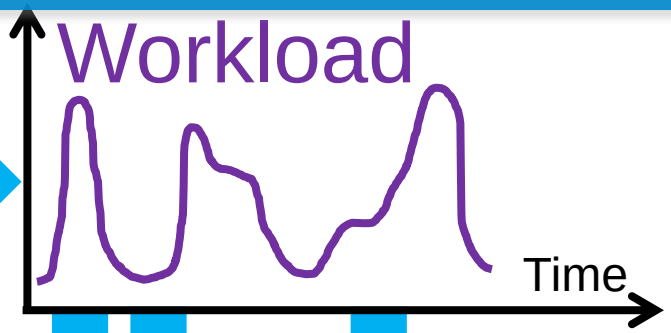

100+ services

Scheduler

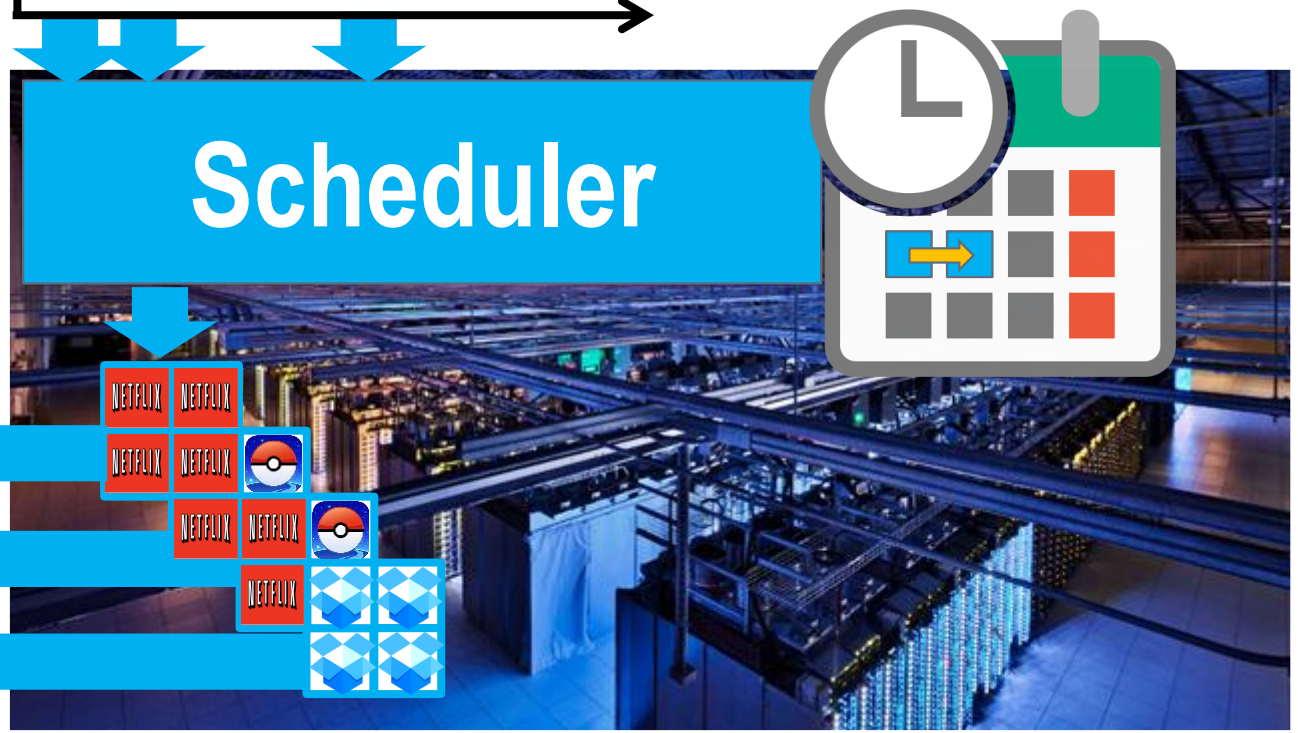

(full-/micro-/nano-) Datacenter

Performance, Dependability, Efficiency

My Research: Massivizing Computer Systems

Is 56% uptime good? 66%? 96%?

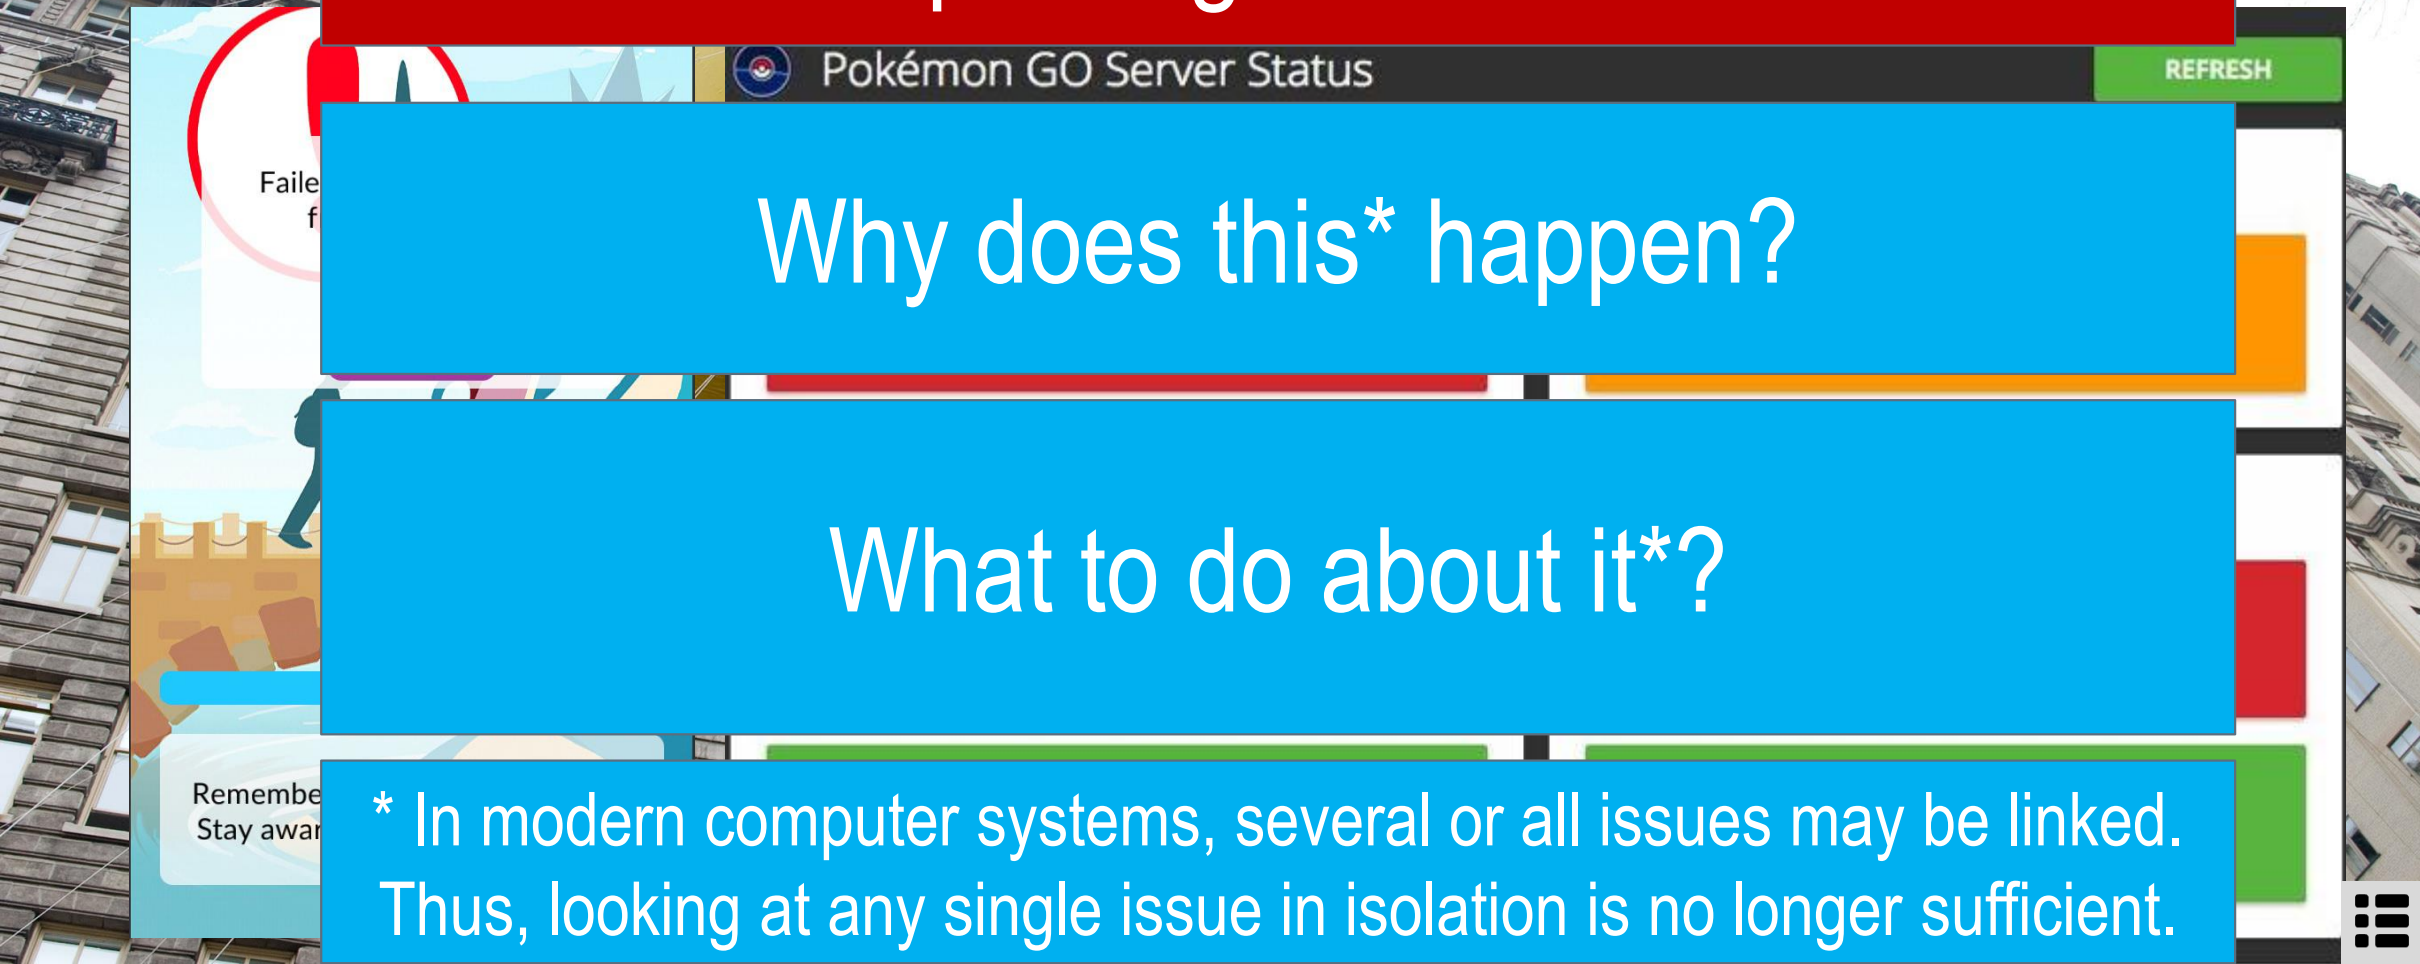

Why does this* happen?

What to do about it*?

* In modern computer systems, several or all issues may be linked.
Thus, looking at any single issue in isolation is no longer sufficient.

THE CRISIS: IN THIS DIGITAL ECONOMY, FEW CAN BE SUCCESSFUL!

THE COMPLEXITY CHALLENGE

1. Ecosystem ≠ Single System/Stack

We Build and Test
Isolated Computer Systems
(or Silos, or Narrow Stacks), Yet
Everything Works in Ecosystems

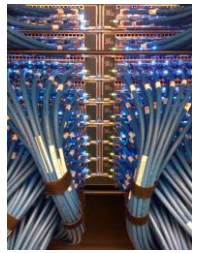

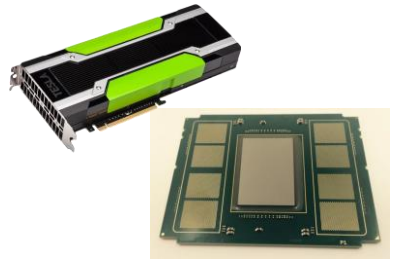

<<1% OF BIG DATA BY MATT TURK (2017)
"SW. IS EATING THE WORLD"

HPC+BIG DATA
CONVERGENCE
"HARDWARE IS THE
NEW SOFTWARE"

THE CRISIS: IN THIS DIGITAL ECONOMY, FEW CAN BE SUCCESSFUL!

PERFORMANCE, DEPENDABILITY, AND OTHER NON-FUNCTIONAL CHALLENGES

2. Non-functional Challenges Not Met

We Cannot Even Maintain the Ecosystems we Have Built (and Tested, and Validated)

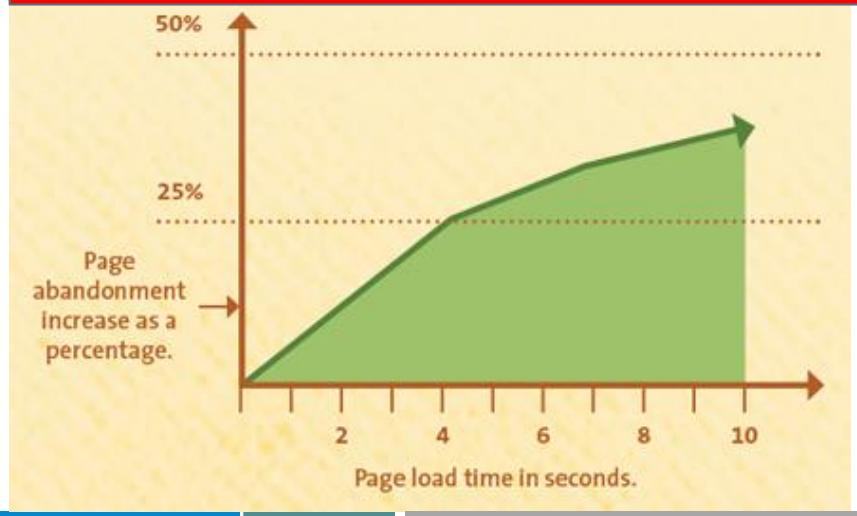

Google goes dark for 2 minutes, kills 40% of world's net traffic www.theregister.co.uk/2013/08/17/google_outage/

Systemwide outage knocks every service offline

Sources: https://www.fastcompany.com/1825005/how-one-second-could-cost-amazon-16-billion-sales

THE CRISIS: IN THIS DIGITAL ECONOMY, FEW CAN BE SUCCESSFUL!

THE RESOURCE MANAGEMENT CHALLENGE

Based on Jay Walker's recent TED talk.

3. Need To Be Much More Efficient,

4. Need to Also Be Ethical, and to Educate Our Clients

PSY Gangnam consumed ~500GWh

= more than entire countries* in a year (*41 countries),
= over 50MW of 24/7/365 diesel, 135M liters of oil,
= 100,000 cars running for a year, ...

Source: Ian Bitterlin and Jon Summers, UoL, UK, Jul 2013.
Note: Psy has >3.5 billion views (last update, May 2018).

THE CRISIS: IN THIS DIGITAL ECONOMY, FEW CAN BE SUCCESSFUL!

THE WORKFORCE GAP, IN THE NETHERLANDS & IN EUROPE

5. A Dire Peopleware Problem

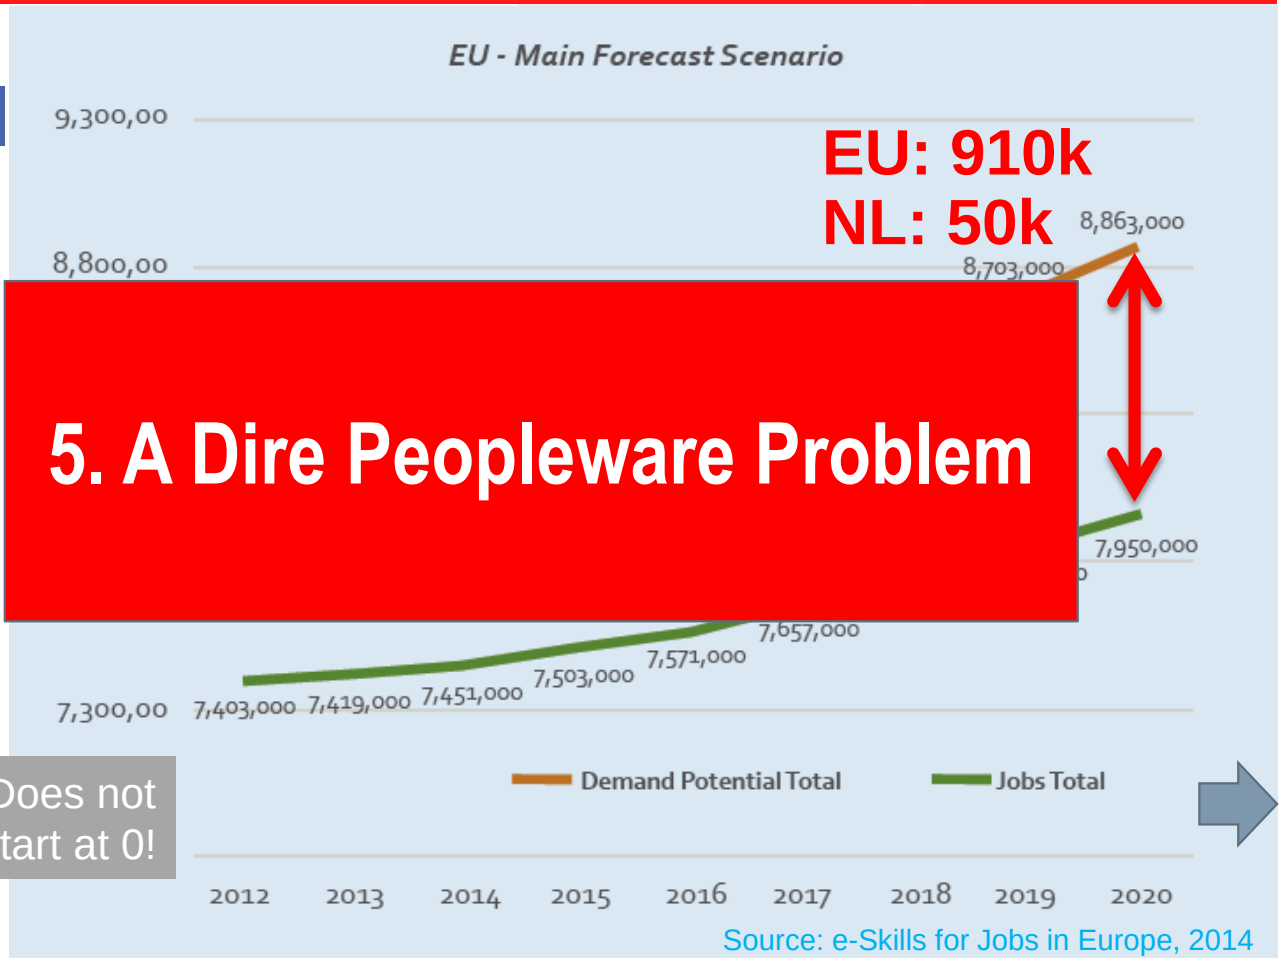

Does not start at 0!

Source: e-Skills for Jobs in Europe, 2014

VU UNIVERSITY AMSTERDAM Faculty of Sciences

THIS IS THE GOLDEN AGE OF DISTRIBUTED COMPUTER SYSTEMS

YET WE ARE IN A CRISIS – 5 CORE CHALLENGES

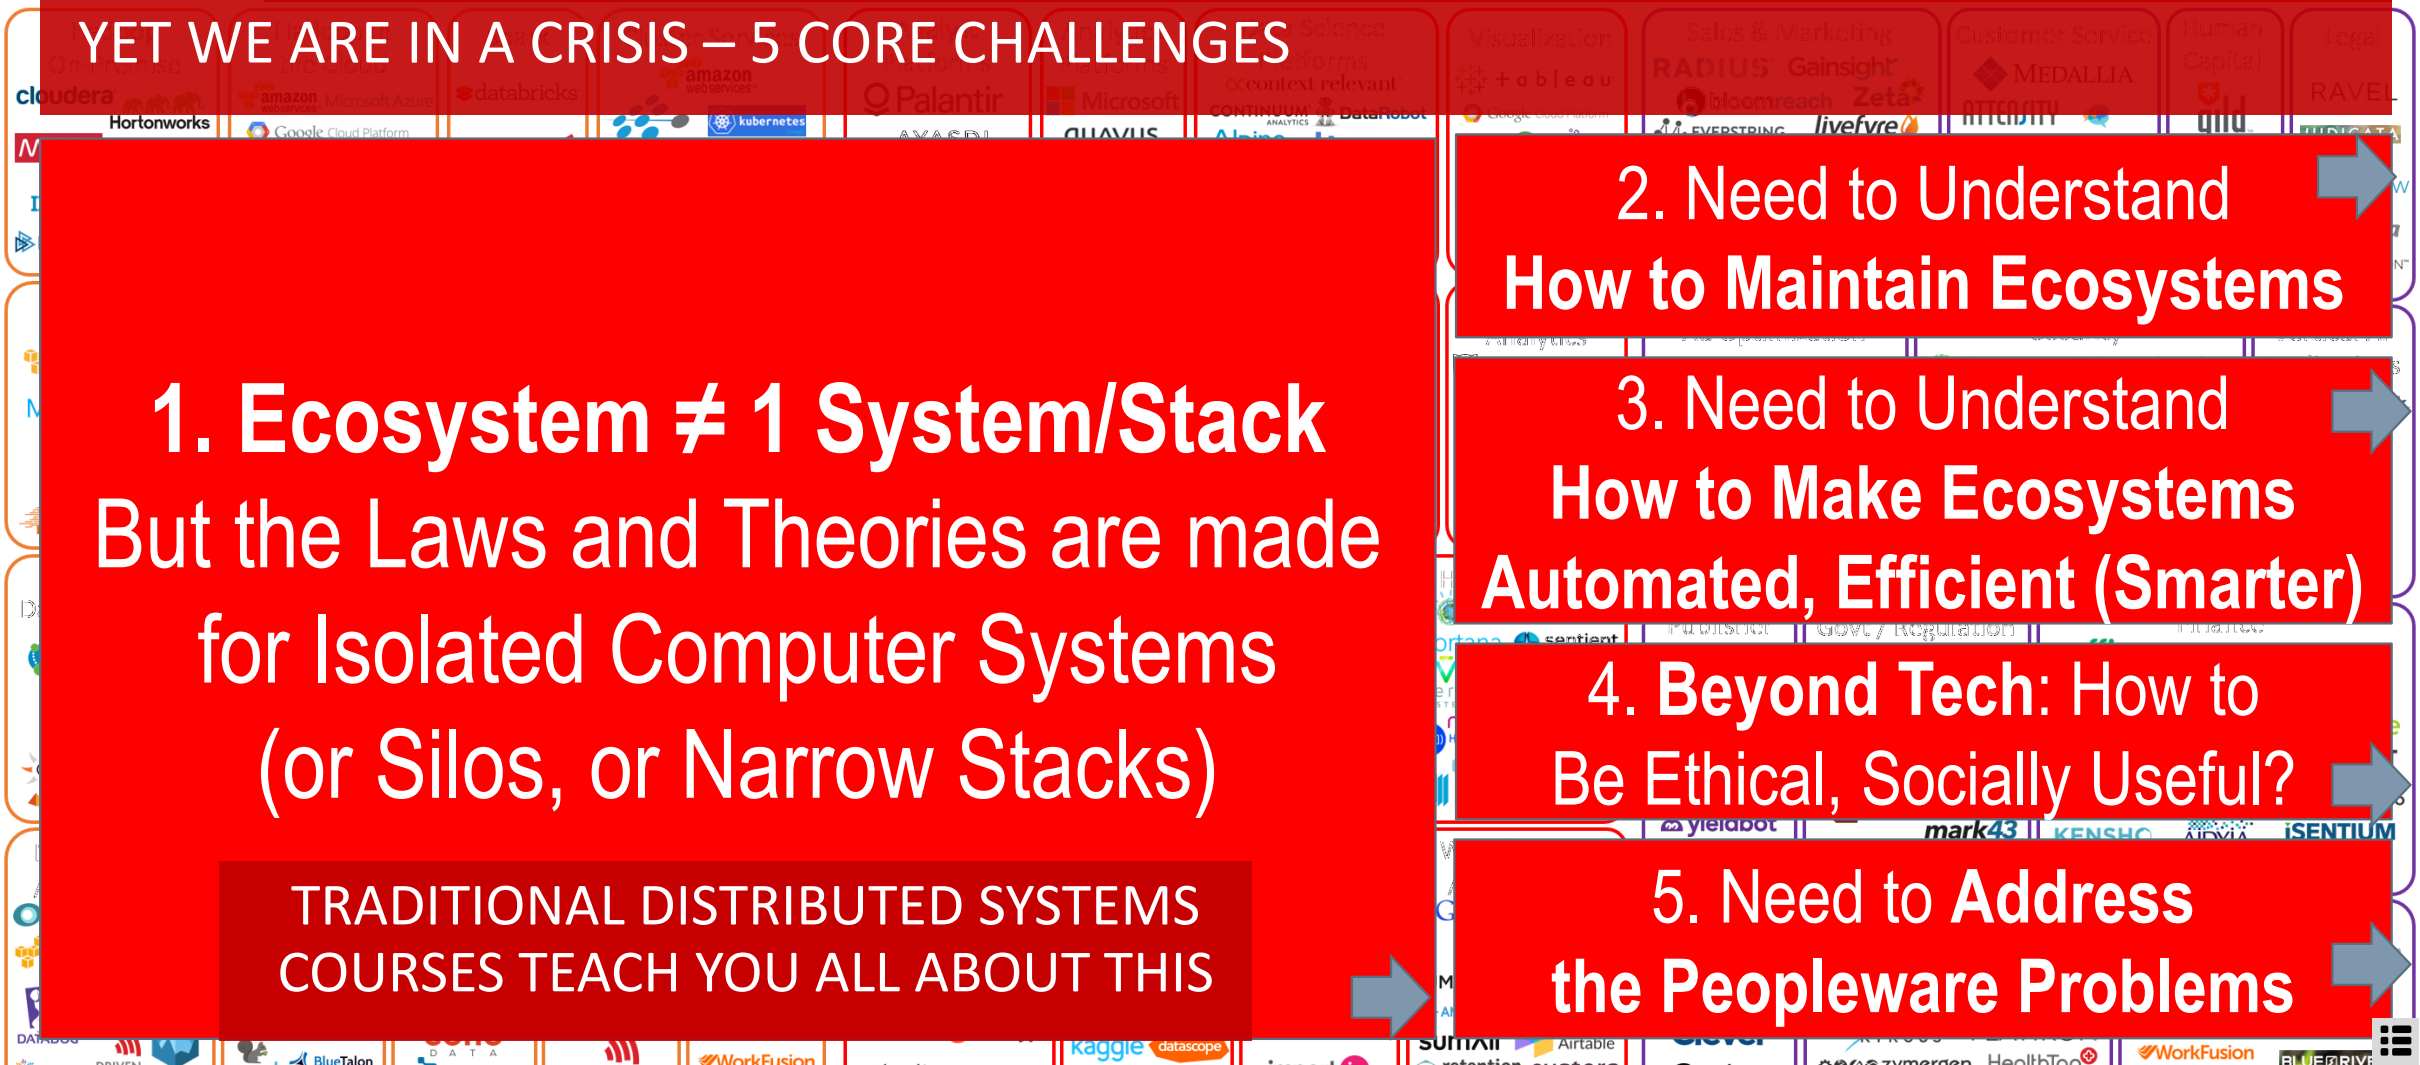

1. Ecosystem ≠ 1 System/Stack

But the Laws and Theories are made for Isolated Computer Systems (or Silos, or Narrow Stacks)

TRADITIONAL DISTRIBUTED SYSTEMS COURSES TEACH YOU ALL ABOUT THIS

2. Need to Understand **How to Maintain Ecosystems**

3. Need to Understand **How to Make Ecosystems Automated, Efficient (Smarter)**

4. **Beyond Tech**: How to Be Ethical, Socially Useful?

5. Need to **Address the Peopleware Problems**

THIS IS THE GOLDEN AGE OF DISTRIBUTED COMPUTER SYSTEMS

YET WE ARE IN A CRISIS

WHICH WE & YOU CAN HELP SOLVE!

Massivizing Computer Systems Tackles The Challenges of Distributed Systems and Ecosystems…

… and Is Relevant, Impactful, and Inspiring for Many Young Scientists and Engineers

OUR WAY FOR DISTRIBUTED SYSTEMS

Massivizing Computer Systems
A Structured Discussion

~60’

~3’ — About Our Team

~30’ — The Golden Age of Distributed Ecosystems … and a Crisis

Interrupts welcome

The main challenges

How we address them: Massivizing Computer Systems

~25’ — Massivizing Computer Systems: Let’s Collaborate

Key for our discussion

What can DistribSys x PerfEng x SwEng do together?

8 ideas for collaboration

~2’ — Take-Home Message

With further reading

THIS IS THE MODERN SCIENCE OF DISTRIBUTED ECOSYSTEMS

MASSIVIZING COMPUTER SYSTEMS IN A NUTSHELL

WHO?	 SCIENTISTS,  ENGINEERS,  DESIGNERS,  MANAGERS, ETC.
WHAT? MAIN GOAL	UNDERSTAND AND CONTROL DISTRIBUTED ECOSYSTEMS, TO TURN THEM INTO EFFICIENT, AUTOMATED UTILITIES
HOW? CENTRAL PARADIGM	ECOSYSTEM OPERATION AND CHARACTERISTICS DERIVE NON-TRIVIALLY FROM ITS SYSTEMS AND USERS (RECURSIVELY)
WHICH APPROACH?	MODERN PROBLEM-SOLVING + DISTRIBUTED (ECO)SYSTEMS

Iosup et al., Massivizing Computer Systems, ICDCS 2018. [Online]

MODERN PROBLEM-SOLVING, MEANINGFUL DISCOVERY

MASSIVIZING COMPUTER SYSTEMS IN A NUTSHELL

science + engineering + design

THE COMPUTER SYSTEMS TRIPLET

[Iosup et al. ICDCS'18]

EXPERIMENTAL METHODS OF DISCOVERY

UNIQUE OPPORTUNITY TO VALIDATE: WE DRINK OUR OWN CHAMPAGNE (*IN VIVO*)!

Our Prototypes (*in phyisico/in vitro*)

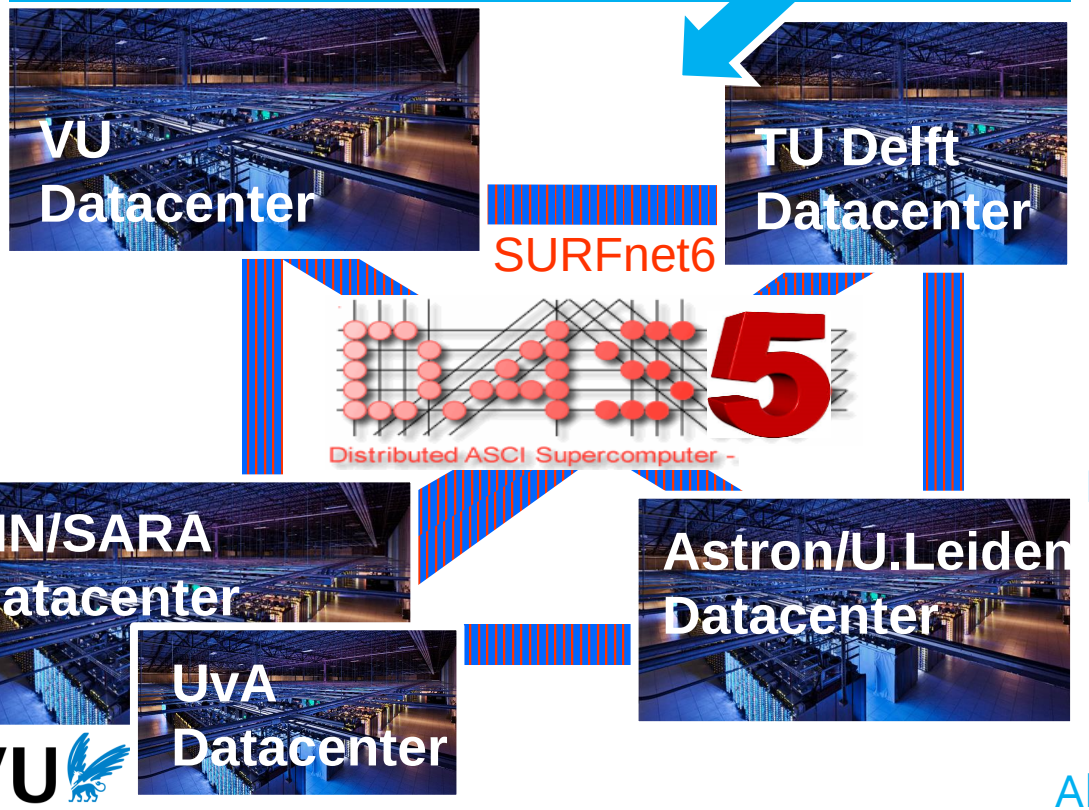

Laurens Versluis

Georgios Andreadis

Fabian Mastebroek

Nacheendra

Maria Voinea

Alexey Ilyushkin

\+ amazon web services™

We also use clouds

\+ OpenDC

And simulators (*in silico*)

NO MORE ARTIFICIAL BOUNDARIES → UNIQUE COLLABORATION

MASSIVIZING COMPUTER SYSTEMS IN A NUTSHELL

Autonomy, Consistency, Composability, *Elasticity*, etc.

Distributed Systems and Ecosystems

Requirements, APIs, DevOps, Architecture and Patterns, etc.

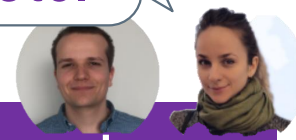

Software Engineering

Performance, *Efficiency*, DevOps, Monitoring, Measurement, etc.

Performance Engineering

Iosup et al., Massivizing Computer Systems, ICDCS 2018. [Online]

~60’

Massivizing Computer Systems
A Structured Discussion

~3’ — About Our Team

~30’ — The Golden Age of Distributed Ecosystems … and a Crisis

Interrupts welcome

The main challenges

How we address them: Massivizing Computer Systems

~25’ — Massivizing Computer Systems: Let’s Collaborate

Key for our discussion

What can DistribSys x PerfEng x SwEng do together?

8 ideas for collaboration

~2’ — Take-Home Message

With further reading

HOW CAN WE COLLABORATE? 8* IDEAS FOR TODAY

DISCUSSION ACROSS DISTRIB. SYS X PERF. ENG. X SW. ENG.

SUGGESTION FOR COLLABORATION

1. > Find together phenomena in ecosystems
2. > Map together artifacts and concepts
3. > Manage together requirement engineering and other DevOp processes
4. > Automate together testing, validation, benchmarking
5. > Localize together faults/issues, identify problems, propose repairs
6. > Process event/data streams using serverless and big data
7. > Manage risks in distributed ecosystems (clouds, big data, edge/IoT, etc.)
8. > Establish the bases of reproducibility, validation, and open science

* 8 is a round number in my world

Iosup et al., Massivizing Computer Systems, ICDCS 2018. [Online]

DISCOVERY = LARGE-SCALE, LONG-TERM STUDY

UNCOVERING THE MYSTERIES OF OUR PHYSICAL UNIVERSE

GEORGE SMOOT
NOBEL PRIZE 2006

SKA FUNDING: 500+ FTE, EUR 1.5B

James Cordes, The Square Kilometer Array, Project Description, 2009 [Online]

The Square Kilometer Array Factsheet, How much will it cost?, 2012 [Online]

Phil Diamond and Rosie Bolton, Life, the Universe & Computing: The story of the SKA Telescope, SC17 Keynote. [Online]

DISCOVERY = LARGE-SCALE, LONG-TERM STUDY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

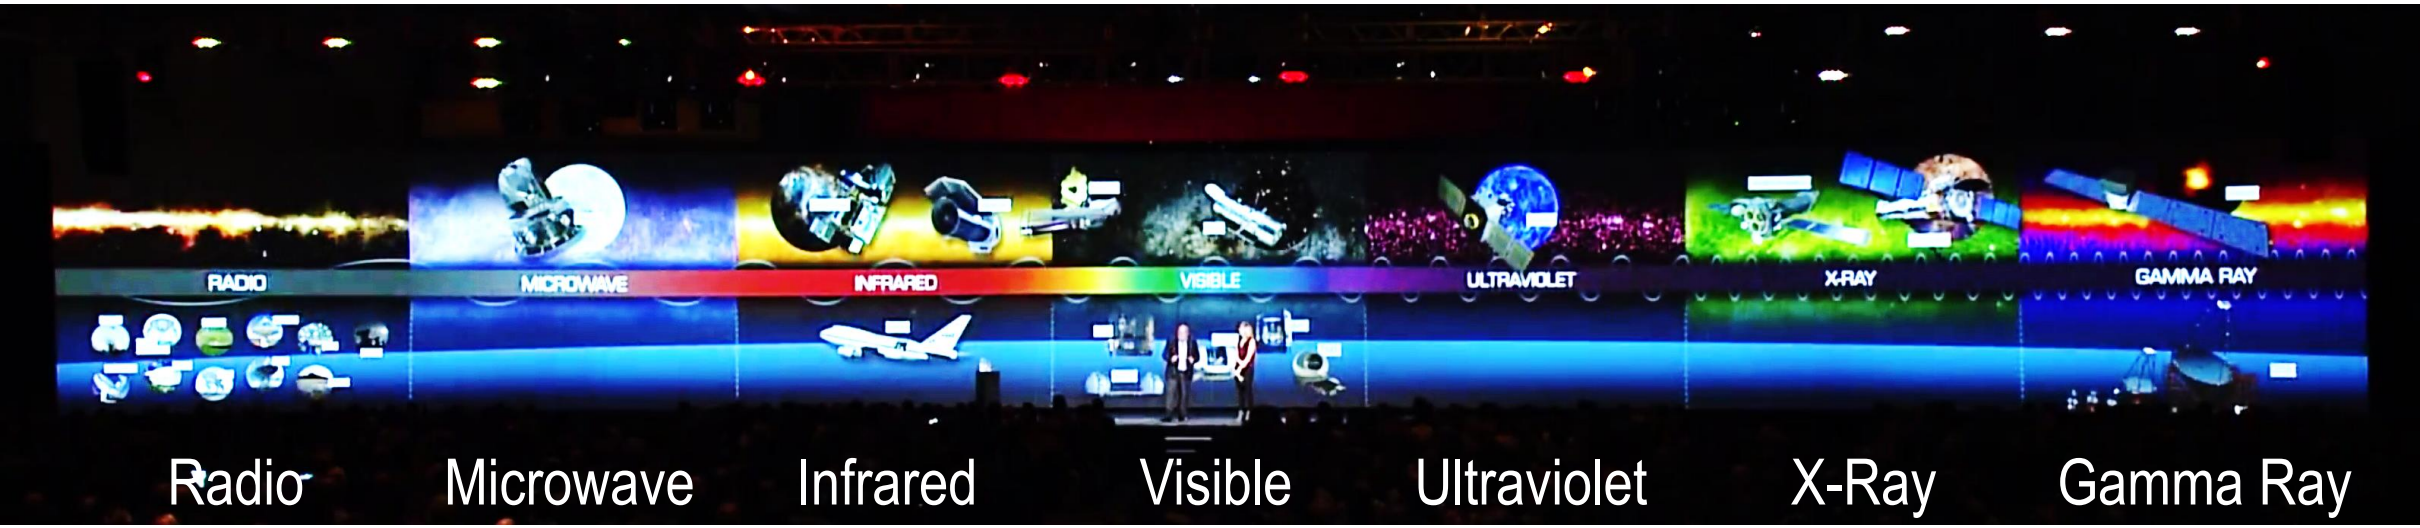

Radio	Microwave	Infrared	Visible	Ultraviolet	X-Ray	Gamma Ray
Cloud, Grid, Edge, Fog, etc.	One aspect: BigData, P2P	Sci.&Eng. Apps+Sys.	Consumer Apps+Sys.	Enterprise Sys.	Systems, Ecosystems	Performance, Availability, etc.
		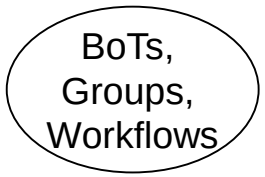				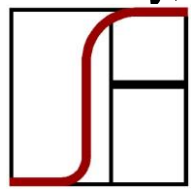
[Iosup et al. FGCS'08]	[Zhang et al. CoNext'10]	[Iosup et al. IEEE IC'11]	[Guo et al. NETGAMES'12]	[Shen et al. CCGRID'15]	[Ghit et al. CCGRID'14]	[Iosup et al. CCGRID'10]

MEANINGFUL DISCOVERY IN DISTRIBUTED ECOSYSTEMS

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

BUT … WHY WOULD YOU NEED TO UNCOVER AN ARTIFICIAL UNIVERSE?! YOU BUILT IT!

Cloud, Grid, Edge, Fog, etc.

[Iosup et al. FGCS'08]

One aspect: BigData, P2P

[Zhang et al. CoNext'10]

Sci.&Eng. Apps+Sys.

[Iosup et al. IEEE IC'11]

Consumer Apps+Sys.

[Guo et al. NETGAMES'12]

Enterprise Sys.

[Shen et al. CCGRID'15]

Systems, Ecosystems

[Ghit et al. CCGRID'14]

Performance, Availability, etc.

[Iosup et al. CCGRID'10]

UNKNOWN PHENOMENA: INTER-, ADAPT-, EXAPTATION

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

@ANDY, ALL: HOW MUCH MORE WE COULD UNDERSTAND BY DISTRIB. SYS. X PERF.ENG. X SW.ENG.?

...	...	...	...	...
BOTS, NOT PARALLEL JOBS	GROUPS NOT RARE, DOMINANT	COMMUNITY FORMATION	SYSTEMIC VARIABILITY	CORRELATED, NOT IID FAILURES
Cloud, Grid, Edge, Fog, etc.	One aspect: BigData, P2P	Sci.&Eng. Apps+Sys.	Consumer Apps+Sys.	Enterprise Sys.	Systems, Ecosystems	Performance, Availability, etc.
		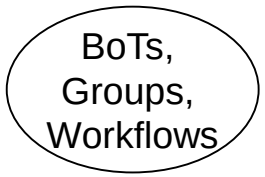				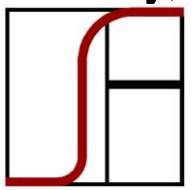
[Iosup et al. FGCS'08]	[Zhang et al. CoNext'10]	[Iosup et al. IEEE IC'11]	[Guo et al. NETGAMES'12]	[Shen et al. CCGRID'15]	[Ghit et al. CCGRID'14]	[Iosup et al. CCGRID'10]

UNKNOWN PHENOMENA: INTER-, ADAPT-, EXAPTATION

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

BUT … IS THERE A SYSTEMATIC WAY TO APPROACH THESE PHENOMENA?

…	…	…	…	…
BOTS, NOT PARALLEL JOBS	GROUPS NOT RARE, DOMINANT	COMMUNITY FORMATION	SYSTEMIC VARIABILITY	CORRELATED, NOT IID FAILURES
Cloud, Grid, Edge, Fog, etc.	One aspect: BigData, P2P	Sci.&Eng. Apps+Sys.	Consumer Apps+Sys.	Enterprise Sys.	Systems, Ecosystems	Performance, Availability, etc.
		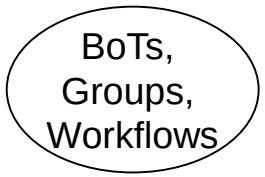				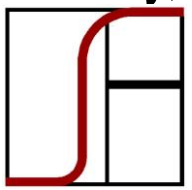
[Iosup et al. FGCS'08]	[Zhang et al. CoNext'10]	[Iosup et al. IEEE IC'11]	[Guo et al. NETGAMES'12]	[Shen et al. CCGRID'15]	[Ghit et al. CCGRID'14]	[Iosup et al. CCGRID'10]

MEANINGFUL DISCOVERY

BUT … IS THERE A SYSTEMATIC WAY TO APPROACH THESE PHENOMENA?

REMEMBER THE COMPLEXITY CHALLENGE?

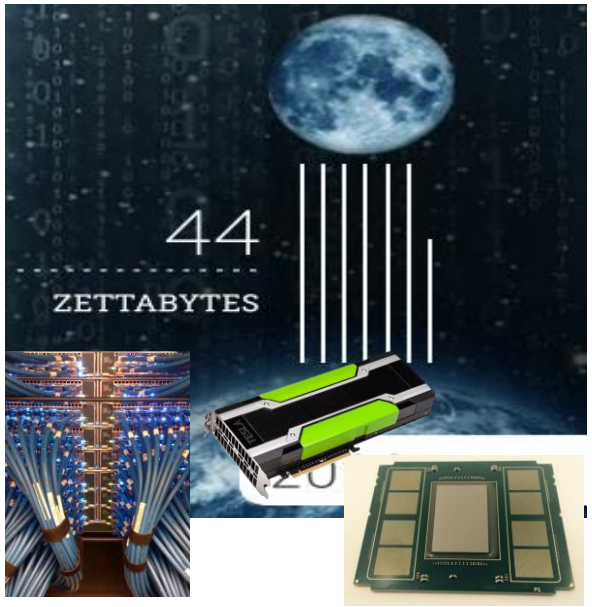

<<1% OF BIG DATA BY MATT TURK (2017)
"SW. IS EATING THE WORLD"

HPC+BIG DATA CONVERGENCE
"HARDWARE IS THE NEW SOFTWARE"

HOW TO MANAGE SYSTEM COMPLEXITY?

THE COMPLEXITY CHALLENGE

Focus on Applications, 5 Core Layers:

5. Development (Front-
4. Runtim (Back-end)
3. Resources
2. Operations Services
1. Infrastructure

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS

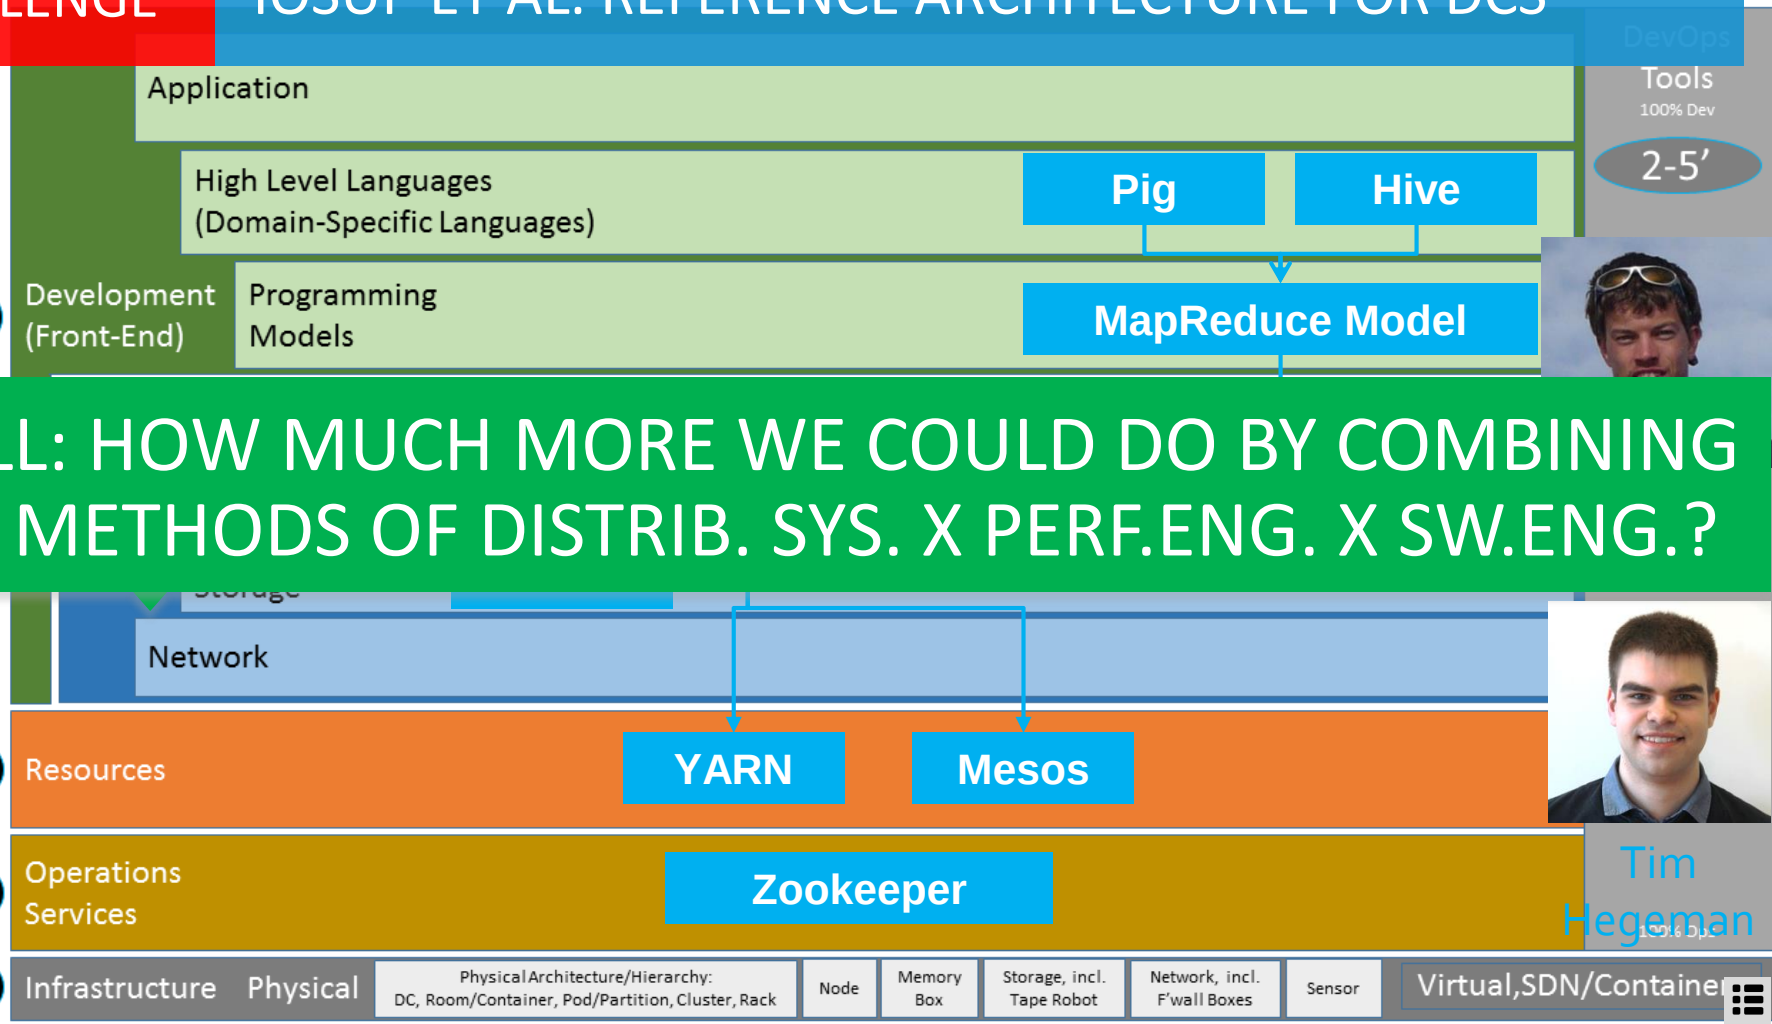

@ARIE, ALL: HOW MUCH MORE WE COULD DO BY COMBINING MAPPING METHODS OF DISTRIB. SYS. X PERF.ENG. X SW.ENG.?

Tim Hegeman

VU [Iosup et al. ICDCS'18]

HOW TO MANAGE SYSTEM COMPLEXITY?

THE COMPLEXITY CHALLENGE

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS

Georgios Andreadis

ANDREADIS ET AL. REFERENCE ARCHITECTURE FOR SCHEDULERS IN DCS

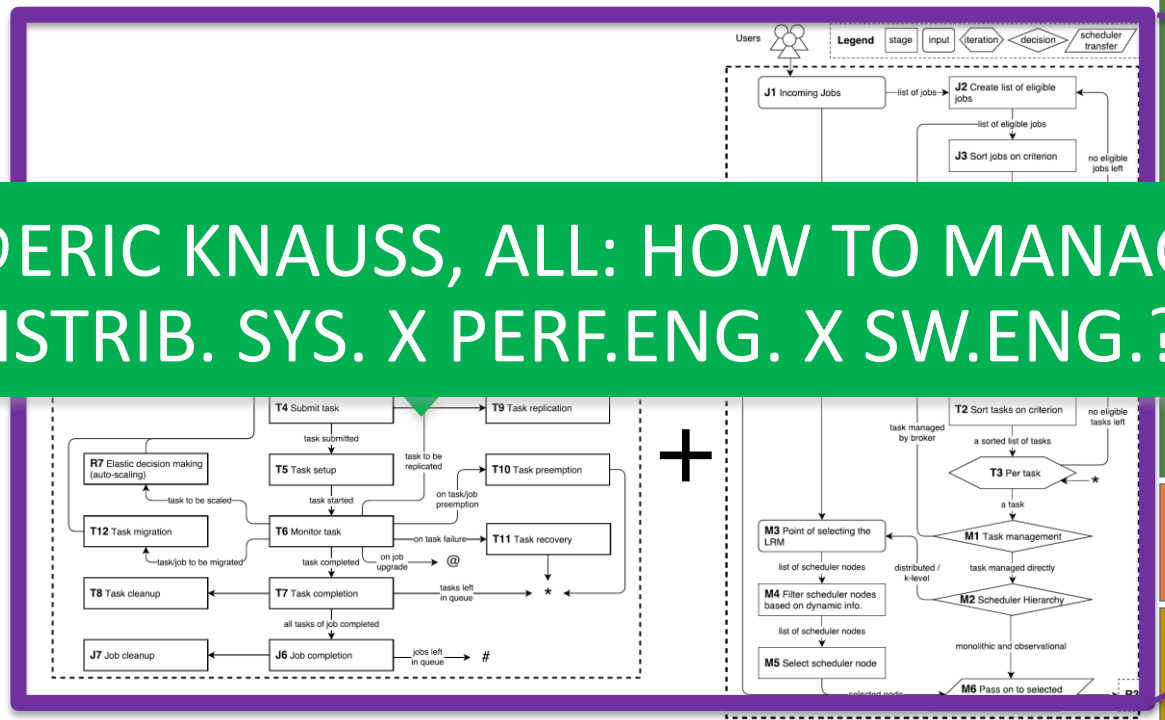

@ERIC KNAUSS, ALL: HOW TO MANAGE REQ. ENG. ACROSS DISTRIB. SYS. X PERF.ENG. X SW.ENG.?

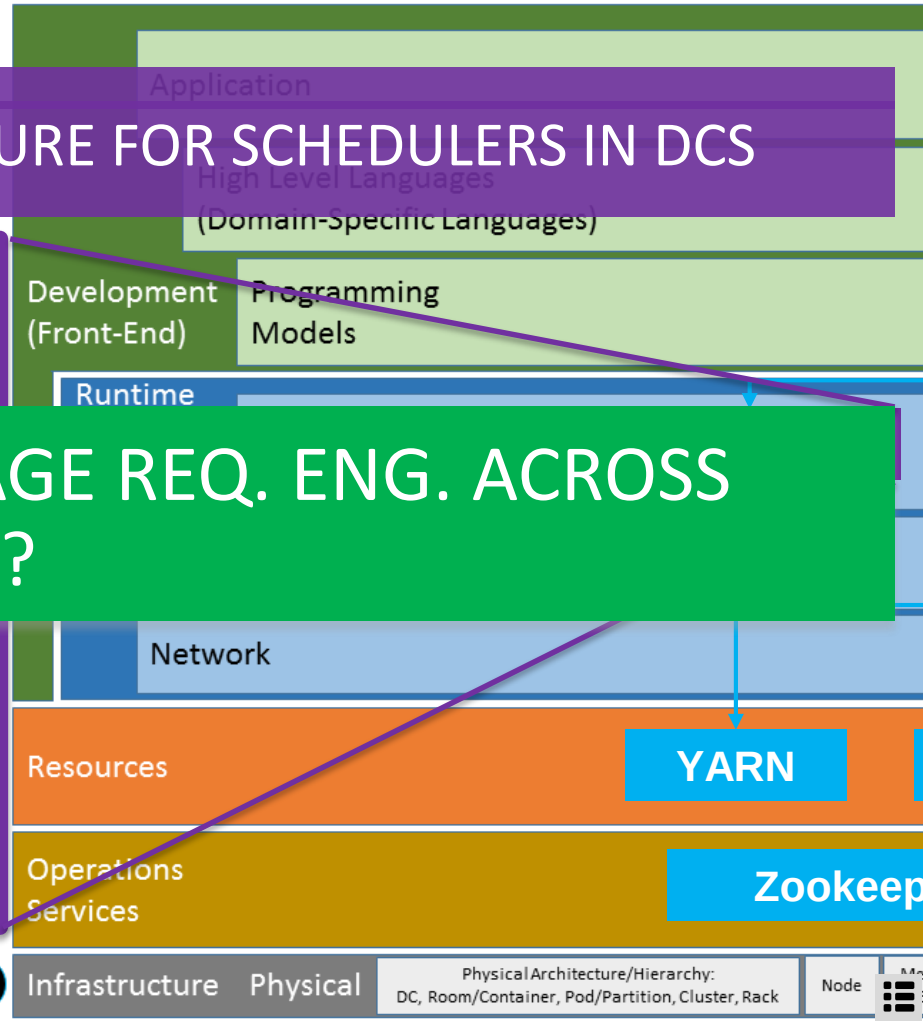

[Andreadis et al. sc'18]

AUTOMATED TESTING FOR DISTRIBUTED ECOSYSTEMS?

ENGINEERING LDBC GRAPHALYTICS: BENCHMARKING LEADING TO DISCOVERY

LDBC The graph & RDF benchmark reference

- Graphalytics:
  - > Ben
  - > Ma
  - > Div
  - > Diverse experiments, representative for practice
  - > Renewal process to keep the workload relevant
  - > Enables comparison of many platforms, community-driven and industrial
  - > Global Competition

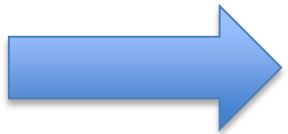

ORACLE intel Georgia Tech IBM HUAWEI CWI UPC

- Community endorsed:

@BENOIT BAUDRY, ALL: MEANINGFUL TESTING, AUTOMATED W/ METHODS ACROSS DISTRIB. SYS. X PERF.ENG. X SW.ENG.?

Performance: orders of magnitude difference due to each of platform, algorithm, dataset, and hardware

VU

[Iosup et al. PVLDB'16] [Guo et al. CCGRID'15] [Guo et al. IPDPS'14]

https://graphalytics.org

LOCALIZATION OF BOTTLENECKS → PERF. ISSUES

ENGINEERING LDBC GRAPHALYTICS: MODELING LEADS TO DISCOVERY

LDBC The graph & RDF benchmark reference

- Graphalytics Grade10:
  - > Au
  - > Au

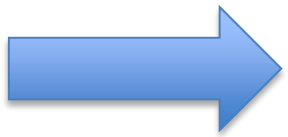

Without Grade10

@CLAIRE LE GOUES, ALL: MEANINGFUL LOCALIZATION OF FAULTS/ISSUES, FOR DISTRIB. SYS. X PERF.ENG. X SW.ENG.?

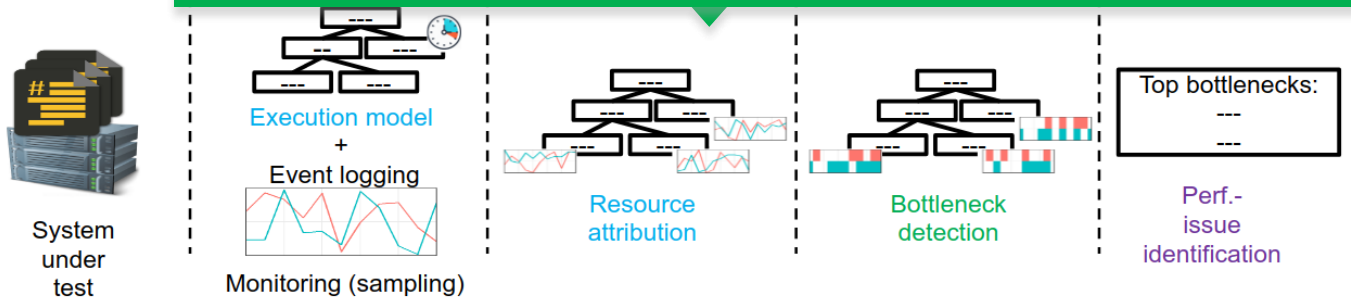

Multi-stage process, works in ecosystem

Always bottleneck
Can explain causes:
+ Message queue full
+ Garbage collector
+ CPU
+ Others

VU

SERVERLESS STREAMING WORKFLOWS

DESIGNING SERVERLESS ARCHITECTURES, APIS, AND SCHEDULERS

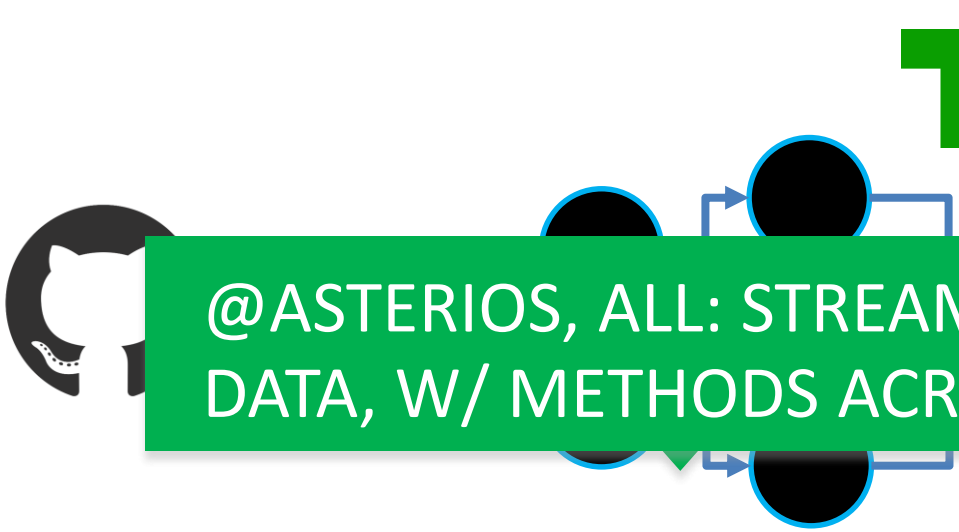

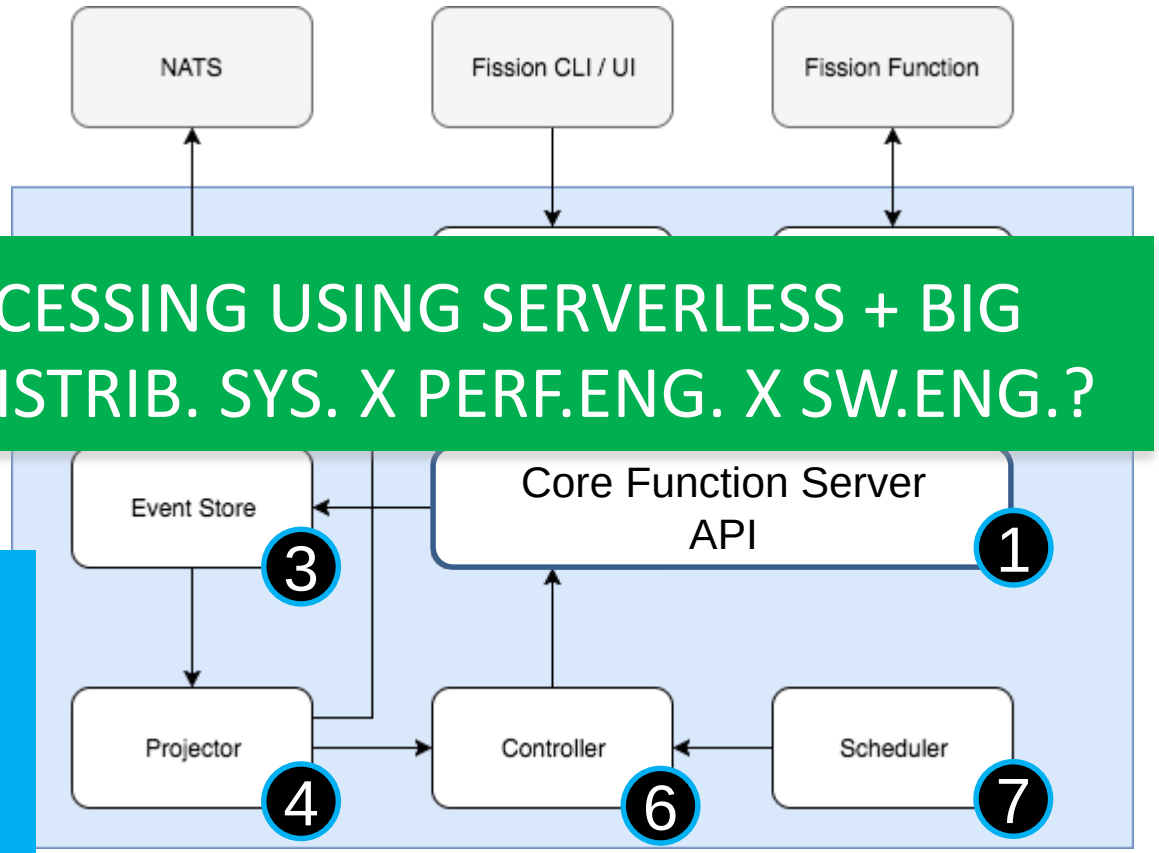

@ASTERIOS, ALL: STREAM PROCESSING USING SERVERLESS + BIG DATA, W/ METHODS ACROSS DISTRIB. SYS. X PERF.ENG. X SW.ENG.?

The first serverless workflow management engine, now part of the Serverless ecosystem at Fission.io

DYNAMIC SCHEDULING TO MANAGE OPERATIONAL RISK

DESIGNING PORTFOLIO SCHEDULERS FOR DATACENTERS, BIG DATA STACKS, ETC.

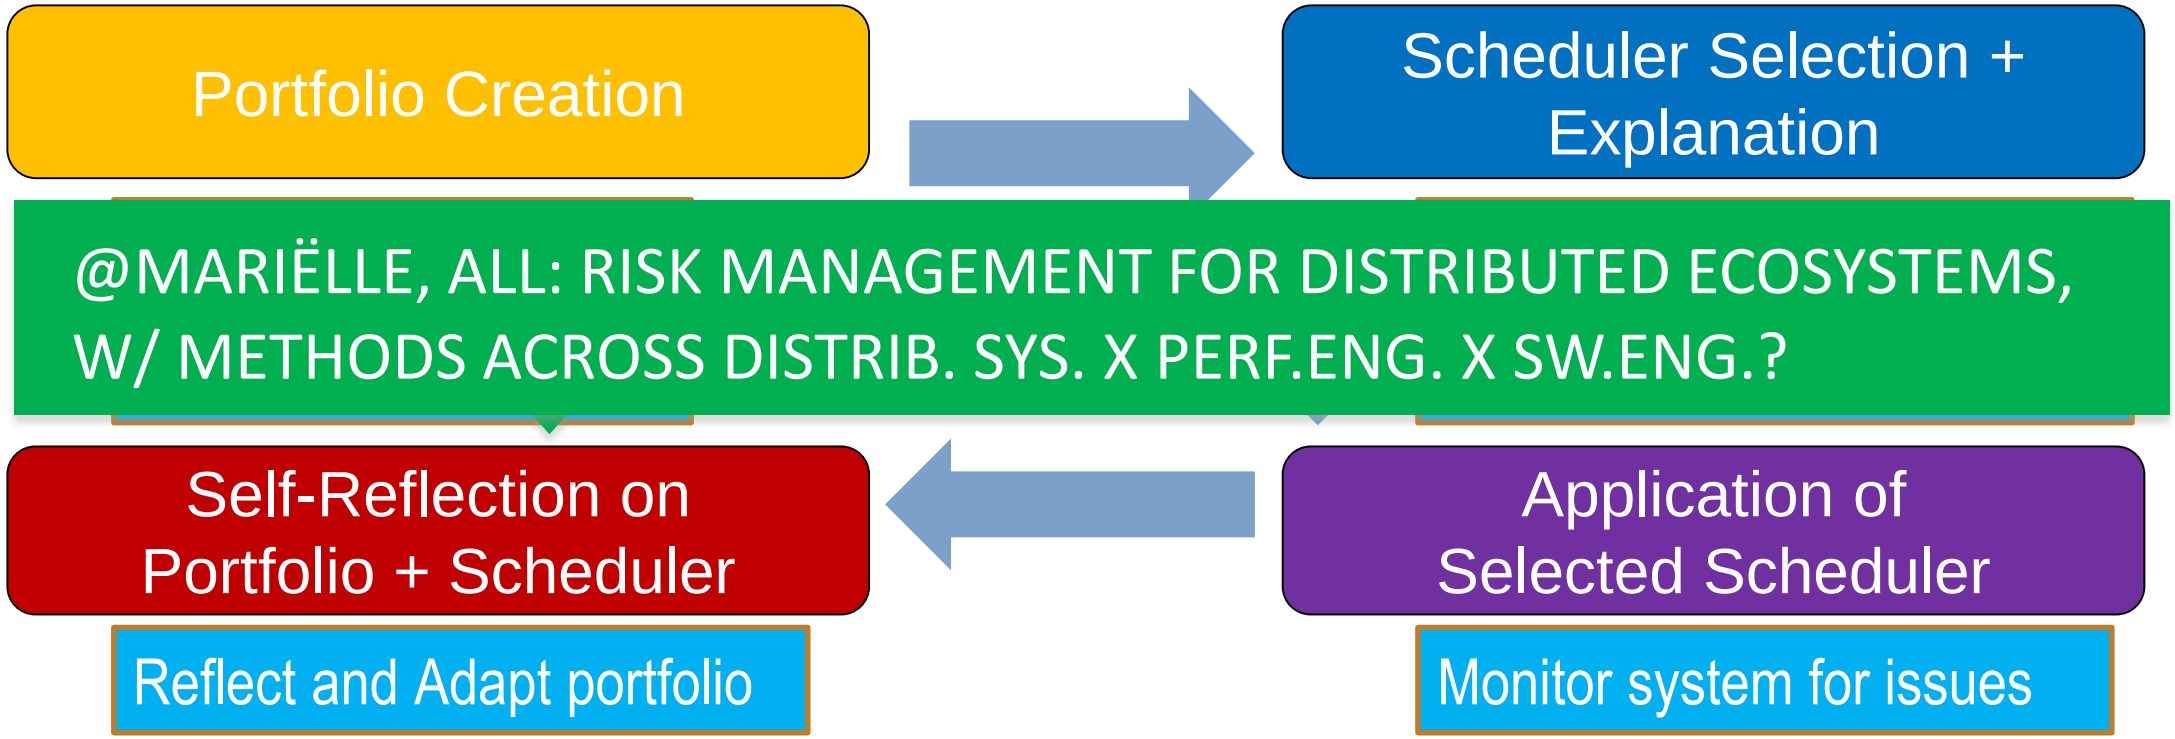

[van Beek et al.

IEEE Computer 2015]

REPRODUCIBILITY AND VALIDATION OF DISCOVERY

A PERENNIALLY TOUGH PROBLEM, IN COMPUTING BUT ALSO IN ALL OTHER SCIENCES

METHODOLOGY	SHARED PRINCIPLES, METHODS, ETC. … BUT WHERE*?!
OP	FREE OPEN SOURCE SOFTWARE AND OPEN ACCESS DATA
REPORTING & DISSEMINATION	PROTOCOL AND STUDY CHECKLISTS, PRE-REGISTRATION OF STUDY AND CONFLICTS-OF-INTEREST … BUT HOW TO START?!
REPRODUCIBILITY	MODERN ECOSYSTEMS ARE NOT STABLE, PREDICTABLE…

@ARIE, ALL: REPRODUCIBILITY, VALIDATION, OPEN SCIENCE, W/ METHODS ACROSS DISTRIB. SYS. X PERF.ENG. X SW.ENG.?

* Conferences do not accept such material… except when they do…

Munafò et al., A manifesto for reproducible science, Nature Human Behaviour, Jan 2017. [Online]

Massivizing Computer Systems
A Structured Discussion

~60’

~3’ — About Our Team

~30’ — The Golden Age of Distributed Ecosystems … and a Crisis

Interrupts welcome

The main challenges

How we address them: Massivizing Computer Systems

~25’ — Massivizing Computer Systems: Let’s Collaborate

Key for our discussion

What can DistribSys x PerfEng x SwEng do together?

8 ideas for collaboration

~2’ — Take-Home Message

With further reading

MASSIVIZING COMPUTER SYSTEMS

= MAKING COMPUTER SYSTEMS SCALABLE, RELIABLE, PERFORMANT, ETC., YET ABLE TO FORM AN EFFICIENT ECOSYSTEM

Many thanks to 200+ collaborators

- Golden Age of Distributed Ecosystems ... Yet a crisis is looming
- Massivizing Computer Systems means modern distributed systems
  - Always Ecosystems
  - Methods to address key challenges in science, design, and engineering
- Much left to do, as we are merely beginning ...
  - You can help!

VU

@Large Research
Massivizing Computer Systems

http://atlarge.science

Contact Me or Our Team

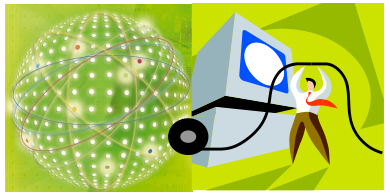

Collaboration or discussion about Massivizing Computer Systems:
Understanding, designing, deploying, tuning, analyzing, benchmarking distributed systems and ecosystems, including cloud computing and big data systems. Other topics in large-scale distributed systems and performance engineering are welcome.

A.Iosup@vu.nl

+31-20 59 89468 (Amsterdam)

@AIosup

https://atlarge-research.com/aiosup/

https://www.linkedin.com/in/aiosup

VU University, Faculty FEW/building W&N, Room P4.14
De Boelelaan 1081, 1081HV Amsterdam,
The Netherlands

VU VRIJE UNIVERSITEIT AMSTERDAM

MASSIVIZING COMPUTER SYSTEMS

FURTHER READING

1. Iosup et al. Massivizing Computer Systems. ICDCS 2018 ← start here
2. Andreadis et al. A Reference Architecture for Datacenter Scheduling, SC18
3. Van Eyk et al. Serverless is More: From PaaS to Present Cloud Computing, IEEE IC Sep/Oct 2018 (in print)
4. Uta et al. Exploring HPC and Big Data Convergence: A Graph Processing Study on Intel Knights Landing, IEEE Cluster 2018

5. Jiang et al. Mirror. CCPE 2018.
6. Ilyushkin et al. Autoscaling for Complex Workflows. TOMPECS 2018.
7. Iosup et al. The OpenDC Vision. ISPDC'17.
8. Iosup et al. Self-Aware Computing Systems book, 2017.
9. Iosup et al. LDBC Graphalytics. PVLDB 2016.
10. Guo et al.: Heterogeneous Graph-Processing. CCGrid 2016.
11. van Beek et al.: IEEE Computer 2015.
12. Jia et al.: TKDD 2015.
13. Ghit et al. SIGMETRICS 2014.
14. Iosup and Epema: IEEE Internet Computing 2011.
15. Iosup et al.: CCGRID 2011.
16. Iosup et al.: IEEE TPDS 2011.

MASSIVIZING COMPUTER SYSTEMS

FURTHER READING II

17. Javadi, Kondo, Iosup, Epema (2013) The Failure Trace Archive: Enabling the comparison of failure measurements and models of distributed systems. J. Parallel Distrib. Comput. 73(8): 1208-1223.
18. Guo and Iosup: The Game Trace Archive. NetGames 2012: 1-6.
19. Iosup et al. (2008) The Grid Workloads Archive. Future Generation Comp. Syst. 24(7): 672-686.
20. Adele Lu Jia et al. (2016) When Game Becomes Life: The Creators and Spectators of Online Game Replays and Live Streaming. TOMCCAP 12(4): 47:1-24.
21. Shen, van Beek, and Iosup: Statistical Characterization of Business-Critical Workloads Hosted in Cloud Datacenters. CCGRID 2015: 465-474.
22. Adele Lu Jia et al. (2015) Socializing by Gaming: Revealing Social Relationships in Multiplayer Online Games. TKDD 10(2): 11:1-29.
23. Iosup et al. (2014) : Analyzing Implicit Social Networks in Multiplayer Online Games. IEEE Internet Computing 18(3): 36-44 (2014).
24. Zhang et al.: Identifying, analyzing, and modeling flashcrowds in BitTorrent. Peer-to-Peer Computing 2011: 240-249.
25. Yigitbasi et al.: Analysis and modeling of time-correlated failures in large-scale distributed systems. GRID 2010: 65-72.
26. Gallet et al.: A Model for Space-Correlated Failures in Large-Scale Distributed Systems. Euro-Par (1) 2010: 88-100.
27. Iosup, Sonmez, and Epema: DGSim: Comparing Grid Resource Management Architectures through Trace-Based Simulation. Euro-Par 2008: 13-25
28. Guo, Hong, Chafi, Iosup, and Epema (2017) Modeling, analysis, and experimental comparison of streaming graph-partitioning policies. J. Parallel Distrib. Comput. 108: 106-21.
29. Guo et al.: Benchmarking graph-processing platforms: a vision. ICPE 2014: 289-292.
30. Uta et al. Elasticity in Graph Analytics? A Benchmarking Framework for Elastic Graph Processing. IEEE Cluster 2018.
31. Heldens, Varbanescu, Iosup: Dynamic Load Balancing for High-Performance Graph Processing on Hybrid CPU-GPU Platforms. IA3@SC 2016: 62-65.
32. Guo et al.: An Empirical Performance Evaluation of GPU-Enabled Graph-Processing Systems. CCGRID 2015: 423-432.
33. Herbst et al. (2016) Ready for Rain? A View from SPEC Research on the Future of Cloud Metrics. CoRR abs/1604.03470 (2016). (in print in TOMPECS)
34. Deng, Song, Ren, and Iosup: Exploring portfolio scheduling for long-term execution of scientific workloads in IaaS clouds. SC 2013: 55:1-55:12.
35. Shen, Deng, Iosup, and Epema: Scheduling Jobs in the Cloud Using On-Demand and Reserved Instances. Euro-Par 2013: 242-254.

© 2017 Alexandru Iosup. All rights reserved.