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

VU VRIJE UNIVERSITEIT AMSTERDAM

ICDCS 2018, Vienna, Austria
Kudos to organizing team!

@Alosup

Prof. dr. ir. Alexandru Iosup

THIS IS THE GOLDEN AGE OF DISTRIBUTED ECOSYSTEMS

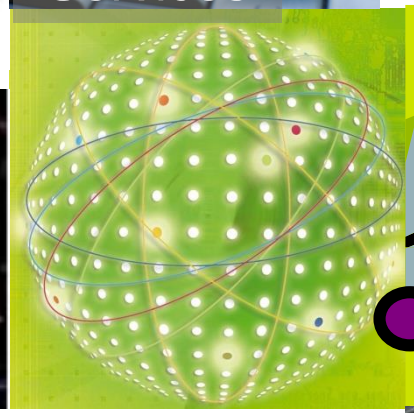

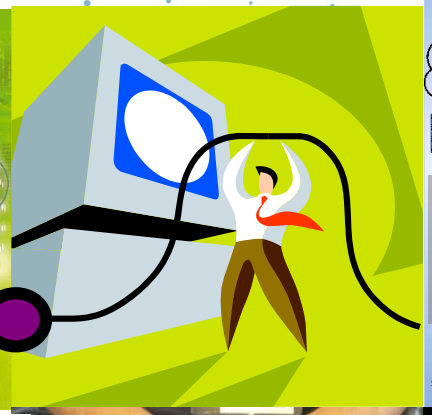

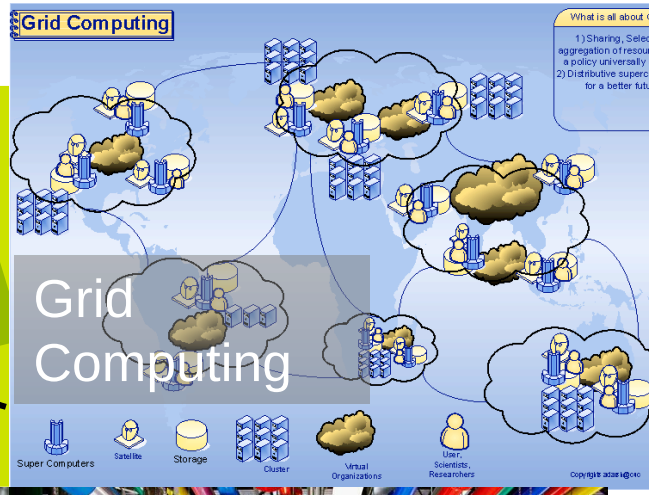

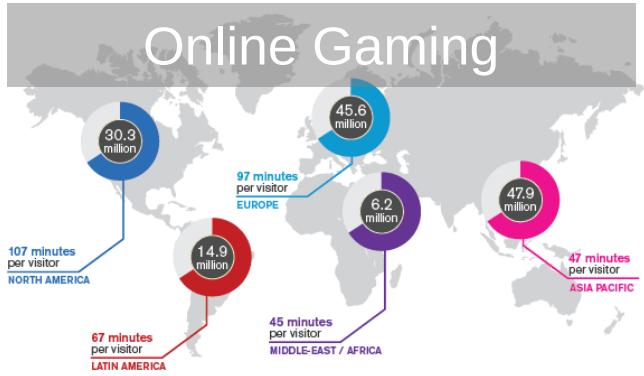

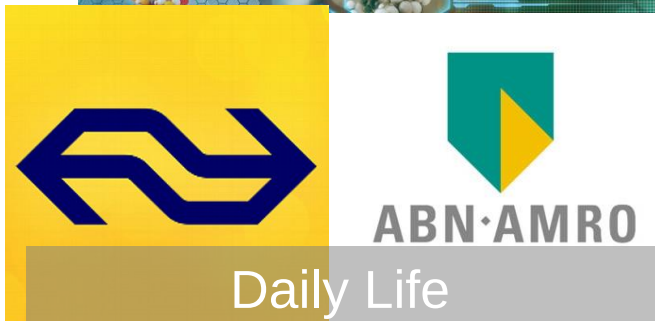

A TYPICAL SYSTEM

DE OLDEHOVE TOWER (16TH CEN.)

> Designed for specific purpose
> Should fit ecosystem needs
> Should not fail for a long time

> Engineering fails in 1529
> Designer tries until death, 1531
> Successor kills the project, 1532
> Project revived repeatedly

> City symbol for 500 years

THIS IS THE GOLDEN AGE OF DISTRIBUTED COMPUTER SYSTEMS

YET WE ARE IN A CRISIS – 5 CORE PROBLEMS RELATED TO ECOSYSTEMS

1. The Current Laws and Theories Are Built and Tested 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: Need to Also Be Ethical

5. Need to Address the Peopleware Problems

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 DISTRIBUTED SYSTEMS AND PROBLEM-SOLVING

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

MEANINGFUL DISCOVERY

science + engineering + design

MANY THANKS TO
180+ CO-AUTHORS

THE COMPUTER SYSTEMS TRIPLET

[Iosup et al. ICDCS'18]

MEANINGFUL DISCOVERY

science + engineering + design

[Iosup et al. ICDCS'18]

MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR 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]

MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

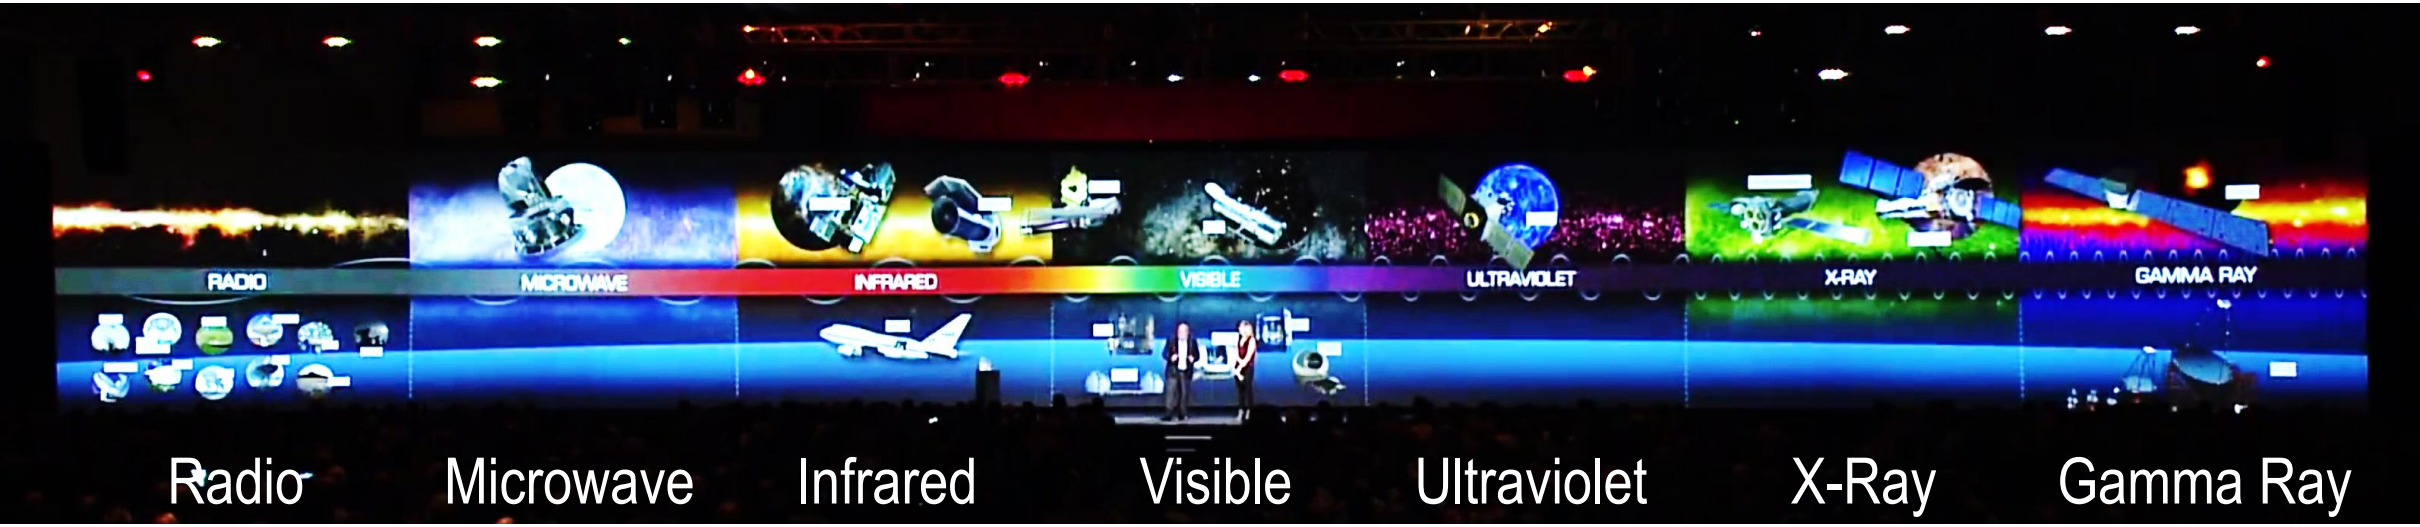

Radio	Microwave	Infrared	Visible	Ultraviolet	X-Ray	Gamma Ray
Cloud, Grid, Edge, Fog, etc.	Big Data	Sci.&Eng. Apps	Consumer Apps	Enterprise Apps	Systems, Ecosystems	Performance, Security, etc.

Domain-Specific Understanding

MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

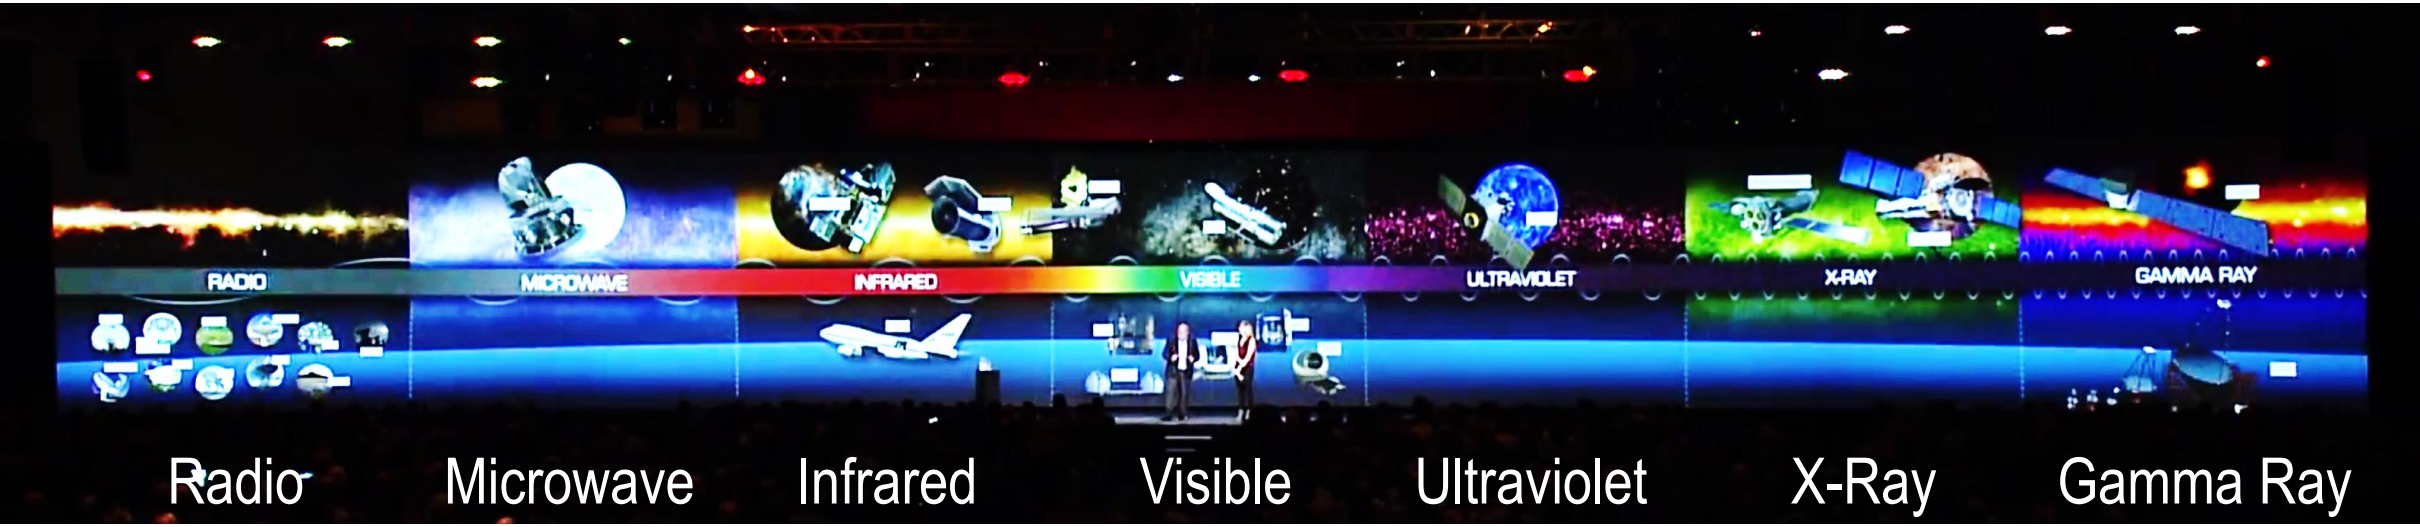

Radio	Microwave	Infrared	Visible	Ultraviolet	X-Ray	Gamma Ray
Cloud, Grid, Edge, Fog, etc.	One aspect: BigData, P2P	Sci.&Eng. Apps+Sys.	Consumer Apps+Sys.	Enterprise Sys.	Systems, Ecosystems	Performance, Availability, etc.
		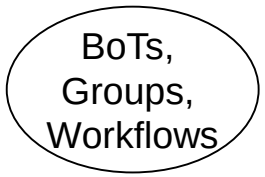				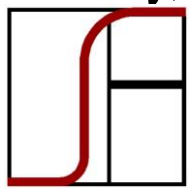
[Iosup et al. FGCS'08]	[Zhang et al. CoNext'10]	[Iosup et al. IEEE IC'11]	[Guo et al. NETGAMES'12]	[Shen et al. CCGRID'15]	[Ghit et al. CCGRID'14]	[Iosup et al. CCGRID'10]

MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

FOUND MANY UNFORESEEN PHENOMENA: INTERACTION, ADAPTATION, EXAPTATION, …

…	…	…	…	…
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?

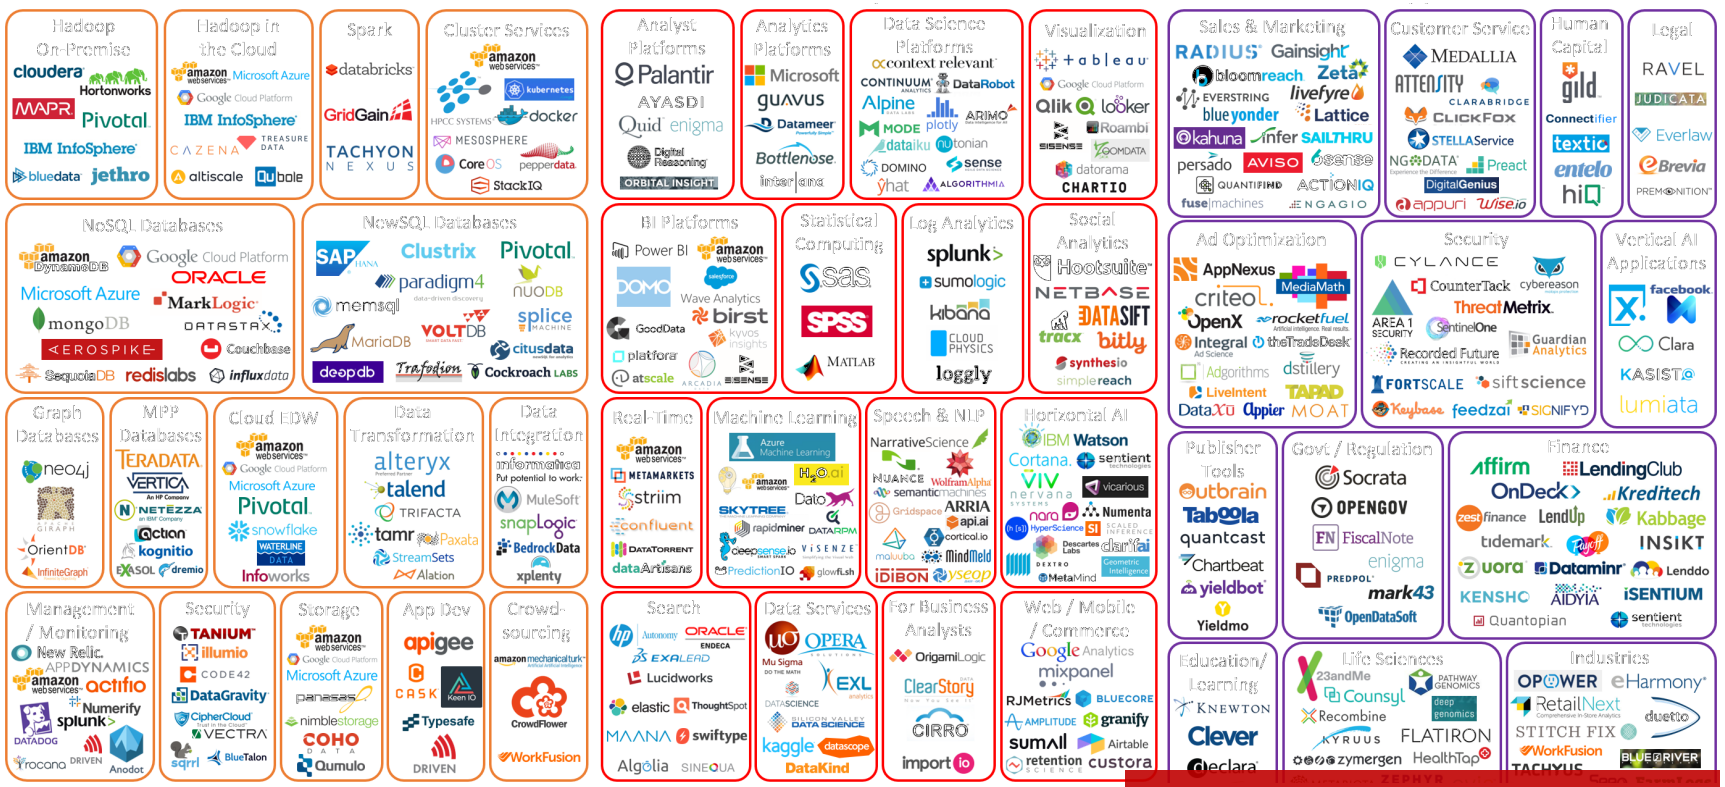

<<1% OF BIG DATA
BY MATT TURK (2017)

MEANINGFUL DISCOVERY

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

International Human Genome Sequencing Consortium, Initial sequencing and analysis of the human genome, Nature 409, Feb 2011. [Online]

Julie Gould, The Impact of the Human Genome Project, Naturejobs blog, 2015. [Online]

MEANINGFUL DISCOVERY

THE COMPLEXITY CHALLENGE

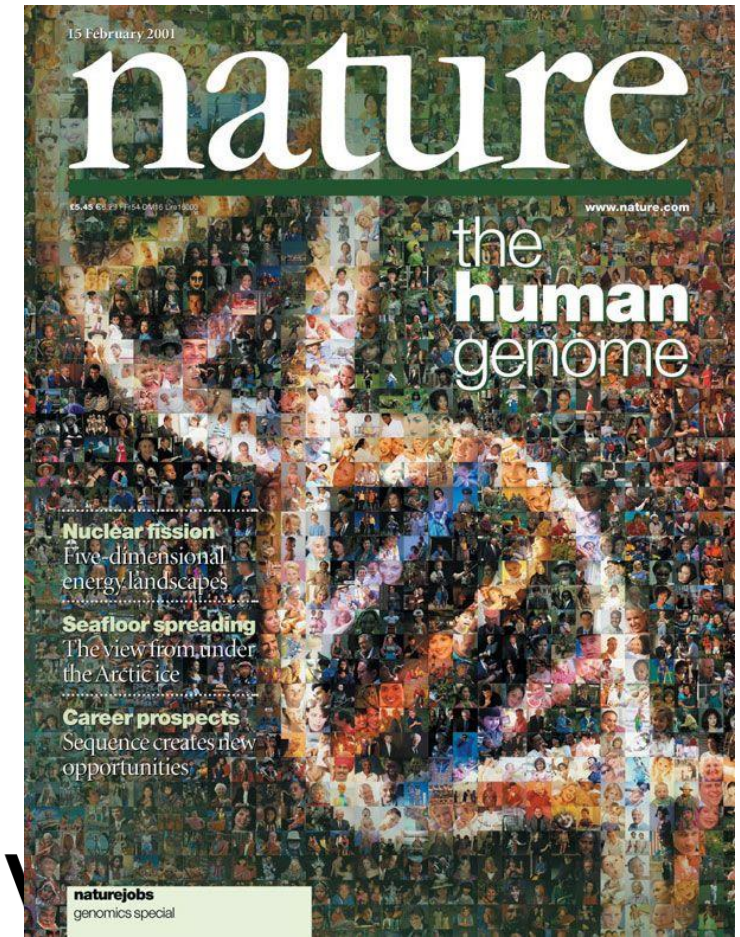

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS

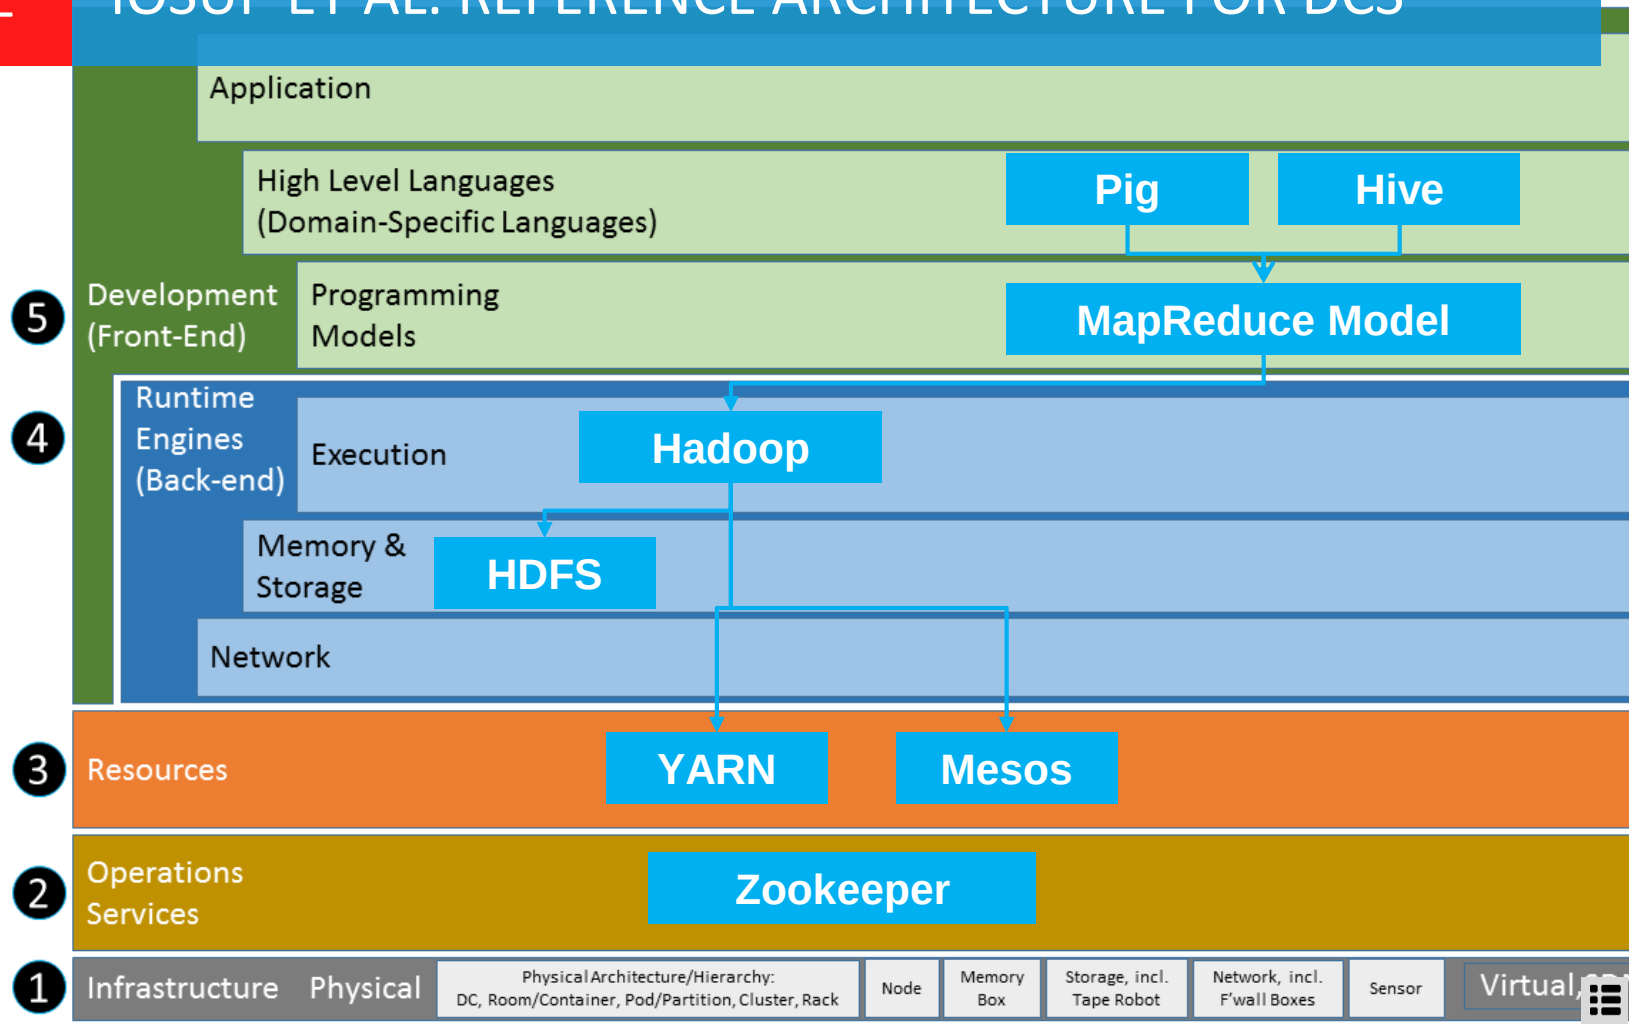

MEANINGFUL DISCOVERY

THE COMPLEXITY CHALLENGE

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS

Georgios Andreadis

ANDREADIS ET AL. REFERENCE ARCHITECTURE FOR SCHEDULERS IN DCS

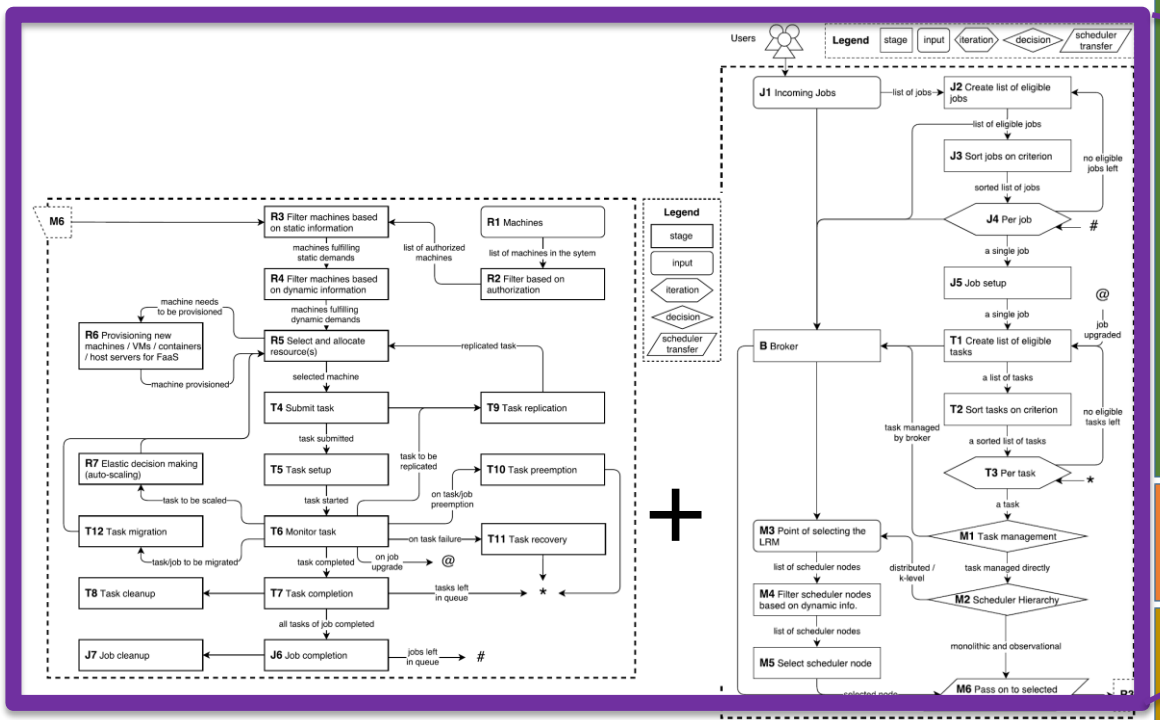

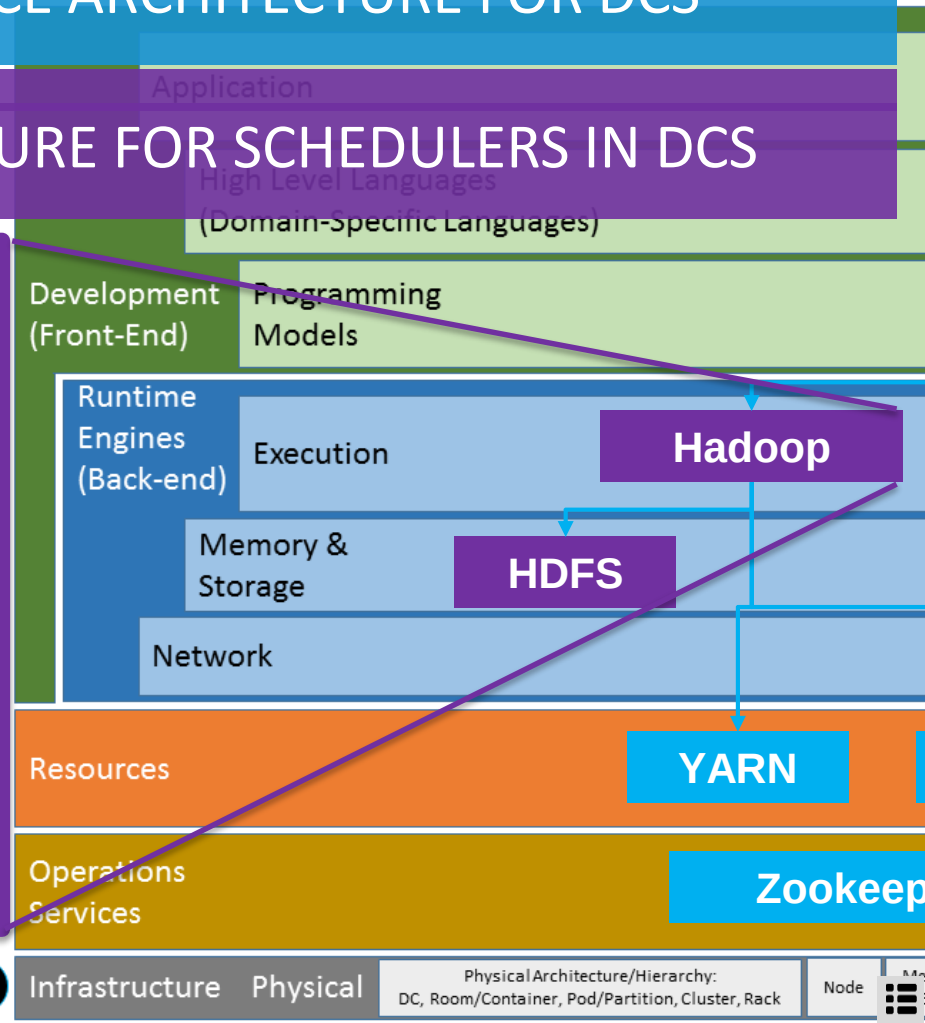

[Andreadis et al. SC'18]

MEANINGFUL DISCOVERY

science + engineering + design

[Iosup et al. ICDCS'18]

MEANINGFUL DISCOVERY

ENGINEERING LDBC GRAPHALYTICS: BENCHMARKING LEADING TO SCIENCE

LDBC The graph & RDF benchmark reference

- Graphalytics:
  - > Benchmark
  - > Many classes of algorithms used in practice
  - > Diverse real and synthetic datasets
  - > Diverse experiments, representative for practice
  - > Renewal process to keep the workload relevant
  - > Enables comparison of many platforms, community-driven and industrial
  - > Global Competition

Wing Lung Ngai

Tim Hegeman

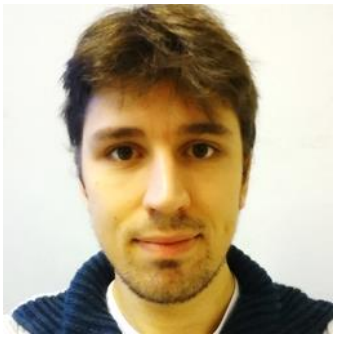

Stijn Heldens

Alex Uță

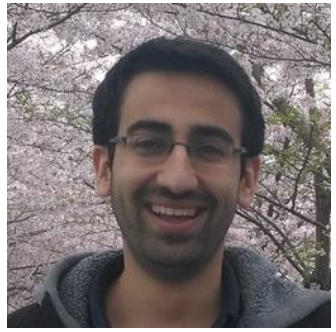

Ahmed Musaafir

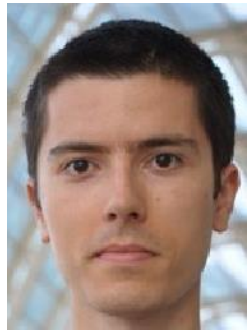

Mihai Capotă

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

MEANINGFUL DISCOVERY

ENGINEERING LDBC GRAPHALYTICS: BENCHMARKING LEADING TO SCIENCE

LDBC The graph & RDF benchmark reference

- Graphalytics:
  - > Benchmark
  - > Many classes of algorithms used in practice
  - > Diverse real and synthetic datasets
  - > Diverse experiments, representative for practice
  - > Renewal process to keep the workload relevant
  - > Enables comparison of many platforms, community-driven and industrial
  - > Global Competition

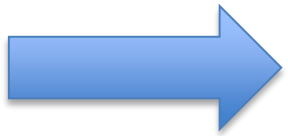

ORACLE intel Georgia Tech IBM HUAWEI CWI UPC

- Community endorsed:

graphalytics.org

- Surprising findings:

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

- Triggered new research

VU

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

MEANINGFUL DISCOVERY

science + engineering + design

[Iosup et al. ICDCS'18]

MEANINGFUL DISCOVERY

DESIGNING SERVERLESS ARCHITECTURES

Erwin
van Eyk

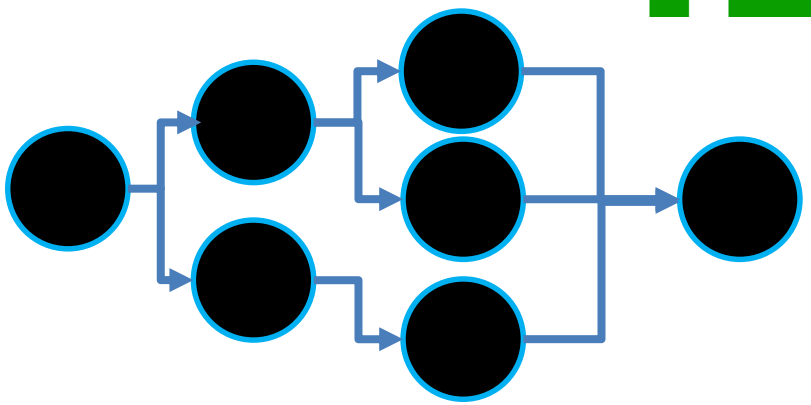

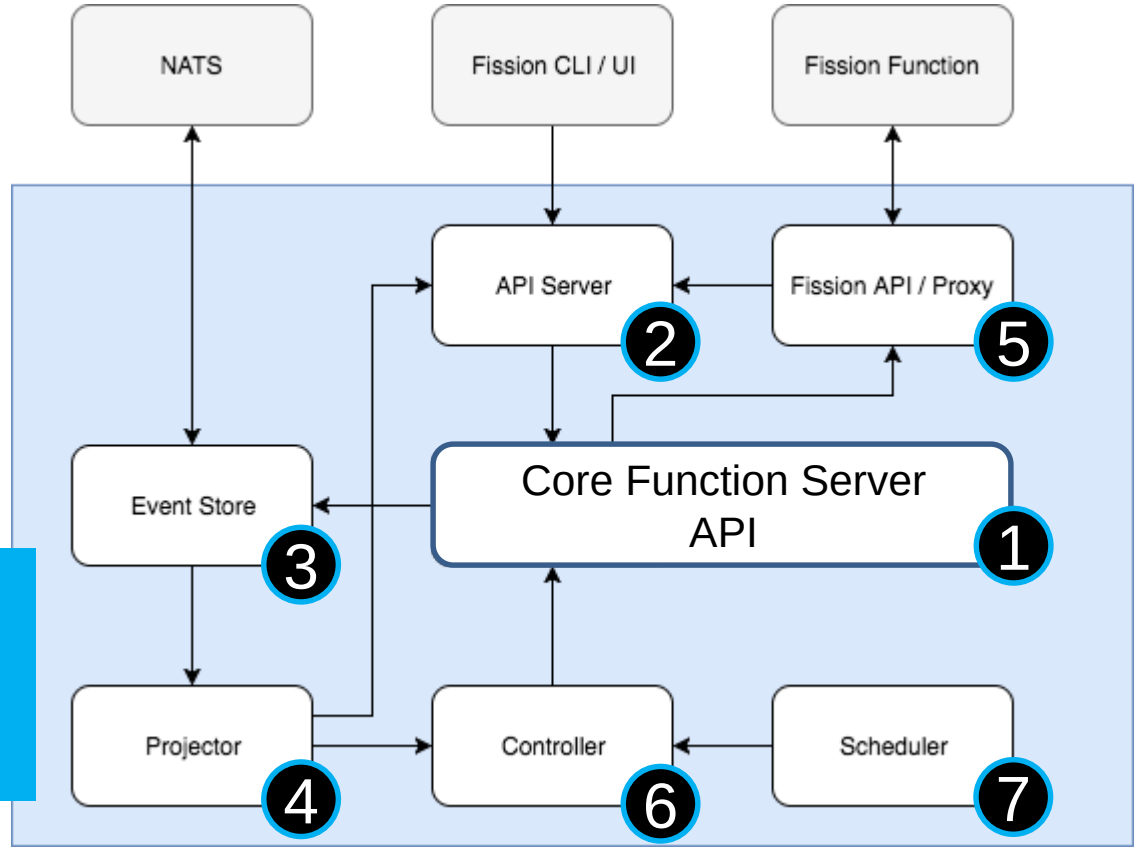

The workflow management engine,
part of the Serverless ecosystem
at Fission.io

https://github.com/fission/fission-workflows/blob/master/Docs/architecture.md

MASSIVIZING COMPUTER SYSTEMS

WHAT ELSE IS IN THE ARTICLE? WHAT ELSE IS IN THE 32-PAGE ARXIV.ORG VERSION?

WHAT?
MAIN GOAL

UNDERSTAND AND CONTROL DISTRIBUTED ECOSYSTEMS, TO TURN THEM INTO EFFICIENT, AUTOMATED UTILITIES

10 Principles + 20 Research Challenges:

- Systems
- Peopleware
- Methodology
- Hapy coincidence: Our principles are a superset of the 7 AI principles at Google

6 Application Domains:

- Datacenter Management
- The Future of Apps
- Generalized Graph Processing
- The Future of Science
- The Future of Online Gaming
- The Future of Banking [in Europe]

Iosup et al., Massivizing Computer Systems, arxiv.org 1802.05465, Feb 2018. [Online]
Sundar Pichai, AI at Google: our principles, Jun 2018. [Online]

MASSIVIZING COMPUTER SYSTEMS

0

WHAT IS ALSO RELATED TO THE ARTICLE? A CALL TO COLLABORATE

1 "a Beckman-like report for Distributed Systems research":

- > Main research challenges
- > Main peopleware challenges
- > Main methodological challenges (new vs. Beckman report on DB research)

2 "a Distributed Systems memex":

- > Datacenter traces
- > Apps traces
- > Operational & performance logs
- > Design logs [*]
- > Implementation logs [*] [not tried / studied before]

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

Iosup et al., Distributed Systems Memex, Dagstuhl Seminar 12472: Is the Future of Data Preservation Cloudy?, 19-21 Nov 2012, Section 5.9. [Online]

MASSIVIZING COMPUTER SYSTEMS

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

Many thanks to 180+ co-authors

- Golden Age of Distributed Ecosystems ... Yet a crisis is looming
- Massivizing Computer Systems means modern distributed systems
  - Think Ecosystems
  - Methods to address key challenges in science, design, and engineering
  - Teaching facilitated by award-winning method ← ask me about this
- Much left to do, as we are merely beginning ...
  - Beckman-like report for Distributed Systems Research
  - The Distributed Systems Memex

spec Research Serverless

VU

@Large Research
Massivizing Computer Systems

http://atlarge.science

MASSIVIZING COMPUTER SYSTEMS

FURTHER READING

1. Iosup et al. Massivizing Computer Systems. ICDCS 2018 (in print)
2. Andreadis et al. A Reference Architecture for Datacenter Scheduling: Design, Validation, and Experiments, SC18 (in print)
3. Van Eyk et al. Serverless is More: From PaaS to Present Cloud Computing, IEEE IC Sep/Oct 2018 (in print)
4. Jiang et al. Mirror: A computation-offloading framework for sophisticated mobile games, CCPE 2018 (in print)
5. Ilyushkin et al. An Experimental Performance Evaluation of Autoscaling Policies for Complex Workflows. TOMPECS 2018.
6. Iosup et al. The OpenDC Vision: Towards Collaborative Datacenter Simulation and Exploration for Everybody. ISPDC'17.
7. Iosup et al. Self-Awareness of Cloud Applications. Self-Aware Computing Systems book, 2017.
8. Iosup et al. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. PVLDB 2016.
9. Guo et al.: Design and Experimental Evaluation of Distributed Heterogeneous Graph-Processing Systems. CCGrid 2016.
10. van Beek et al.: Self-Expressive Management of Business-Critical Workloads in Virtualized Datacenters. IEEE Computer 2015.
11. Jia et al.: Socializing by Gaming: Revealing Social Relationships in Multiplayer Online Games. TKDD 2015.
12. Ghit et al. Balanced resource allocations across multiple dynamic MapReduce clusters. SIGMETRICS 2014.
13. Iosup and Epema: Grid Computing Workloads. IEEE Internet Computing 2011.
14. Iosup et al.: On the Performance Variability of Production Cloud Services. CCGRID 2011.
15. Iosup et al.: Performance Analysis of Cloud Computing Services for Many-Tasks Scientific Computing. IEEE TPDS 2011.

Contact Me or Our Team

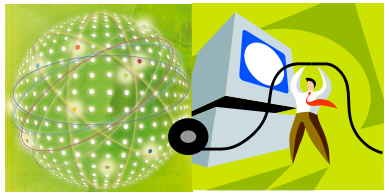

Collaboration or discussion about Massivizing Computer Systems:
Understanding, designing, deploying, tuning, analyzing, benchmarking distributed systems and ecosystems, including cloud computing and big data systems. Other topics in large-scale distributed systems and performance engineering are welcome.

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

WHAT ELSE IS IN THE ARTICLE? WHAT ELSE IS IN THE 32-PAGE ARXIV.ORG VERSION?

SYSTEMS

P1. This is the age of computer ecosystems.

P2. Software-defined everything, but humans can still shape and control the loop.

P3. Non-functional properties are first-class concerns, composable and portable, whose relative importance and target values are dynamic.

P4. Resource Management and Scheduling, and their combination with other capabilities to achieve local and global Self-Awareness, are key to ensure nonfunctional properties at runtime..

P5. Ecosystems are super-distributed.

PEOPLEWARE

P6. People have a fundamental right to learn and to use ICT, and to understand their own use.

P7. Experimenting, creating, and operating ecosystems are professional privileges, granted through provable professional competence and integrity.

P8. We understand and create together a science, practice, and culture of computer ecosystems.

P9. We are aware of the evolution and emergent behavior of computer ecosystems, and control and nurture them. This also requires debate and interdisciplinary expertise.

METHODOLOGICAL

P10. We consider and help develop the ethics of computer ecosystems, and inform and educate all stakeholders about them.

VU