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

THE SCIENCE OF DISTRIBUTED ECOSYSTEMS

@Large Research
Massivizing Computer Systems

http://atlarge.science

VU VRIJE UNIVERSITEIT AMSTERDAM

@Alosup

Prof. dr. ir. Alexandru Iosup

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

- Education:
  - Systems Architecture (BSc)
  - Distributed Systems (MSc)
- Research:
  - Massivizing Computer Systems

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

- Education:
  - > Systems Architecture (BSc)
  - > Distributed Systems (MSc)
- Research:
  - > 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
  - > NL ICT Researcher of the Year
  - > NL Higher-Education Teacher of the Year
  - > NL KNAW Royal Young Academy

VU

NOI RADACINI .nl — 22,000 ROMANIANS IN NL 60% WORKING, 30% STUDENTS (SOME AT VU)

BY CLAUDIA MARCU AND ALEXANDRU IOSUP + MANY VOLUNTEERS

Mihai Netea

http://bit.ly/NR-Mihai

Ana Maria Oprescu

http://bit.ly/NR-AnaMaria

Teodor Cătăniciu

http://bit.ly/NR-Teodor

http://noiradacini.nl

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

- Early vision,
  age 14, in 1994:
  > Computers for everyone
  > Education for everyone
  > The Netherlands

Van Basten.

părinţi. „Jubesc calculatorul şi am câştigat acest concurs tocmai propunând celor mari ca din miliard să cumpere 1.000 de calculatoare pentru a se realiza în şcoli o reţea naţională", spune Alexandru.

ADEVĂRUL

ÎNVĂŢĂMÂNT

"România trebuie să devină o putere"

- consideră un câştigător (ne)obişnuit -

Alexandru Iosup are 14 ani (născut în 12 iunie 1980) şi este elev al Şcolii nr. 19 din Bucureşti. Mic la stat, dar mare la sfat, Alexandru este unul din cei patru câştigători ai concursului „Ce să faci cu un miliard" iniţiat de MTS. Cei 50 de mii de lei primiţi ca premiu al revistei „Modelism" i-a cheltuit cu „cap", luându-şi o minge de fotbal **adevarată** pentru că este microbist convins şi are ca idoli pe Maradona şi Van Basten. În privinţa banilor, copilul are o nedumerire: „Am primit exact 42 de mii lei. De ce s-o fi impozitând şi premiul, nu înţeleg". Nici noi nu pricepem, dar asta e. Oricum, din ce i-a rămas, îşi va luat ceva pentru calculatorul pe care de curând (după premiere) l-a primit ca dar de la părinţi. Jubesc calculatorul şi am câştigat acest concurs tocmai propunând celor mari ca din miliard să cumpere 1.000 de calculatoare pentru a se realiza în şcoli o reţea naţională", spune Alexandru.

vine de la faptul că s-ar simti sărac.

VISION:
WHAT DOES OUR SOCIETY NEED?

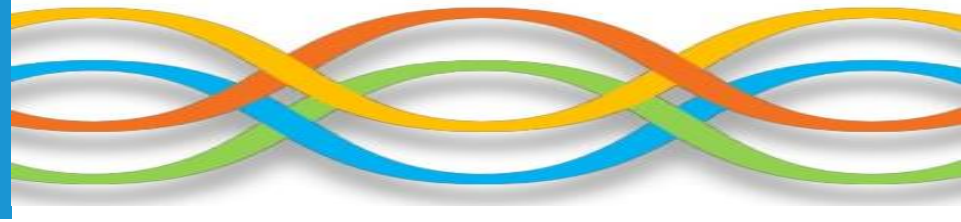

prosperous society

blooming economy

inventive academia

wise governance

VISION: WHAT DOES OUR SOCIETY NEED?

ISN'T THIS ALREADY HAPPENING?

"A world where individuals and human-centered organizations are augmented by an automated, sustainable layer of technology. At the core of this technology is ICT, and at the core of ICT are computer ecosystems, interoperating and performing as utilities and services, under human guidance and control."

- People, good orgs = ICT clients:
  - > Fundamental right to ICT
  - > Understanding
- ICT professionals:
  - > Experiment and create
  - > Operate

- ICT = ecosystems:
  - > Utilities and services
  - > Automated
  - > Efficient
  - > Controlled
  - > Ecosystems
  - > Human-guided

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

THIS IS THE GOLDEN AGE OF DISTRIBUTED ECOSYSTEMS

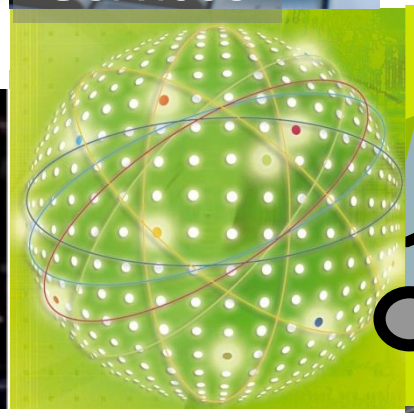

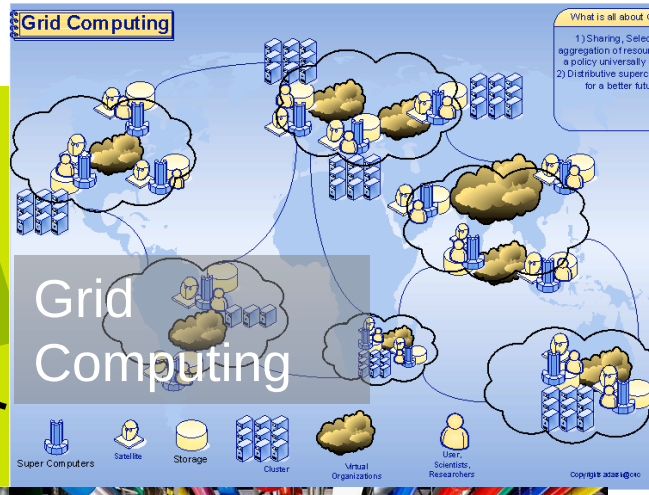

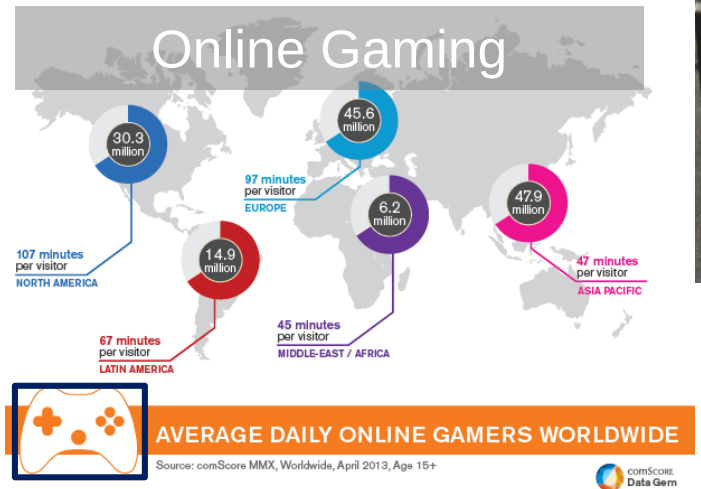

THIS IS THE GOLDEN AGE OF DISTRIBUTED COMPUTER SYSTEMS

Here is how this works…

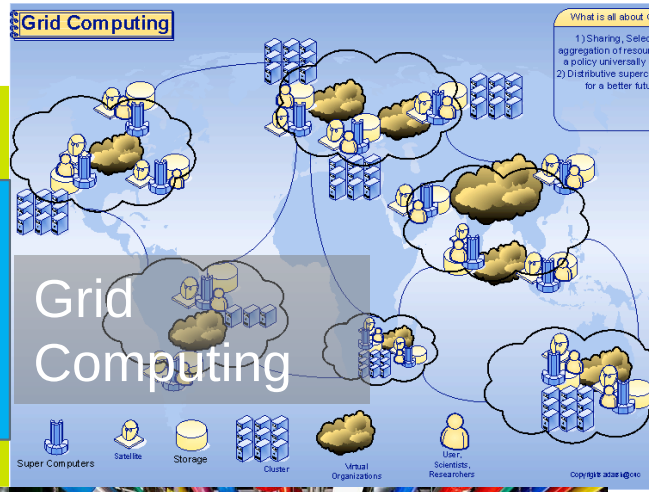

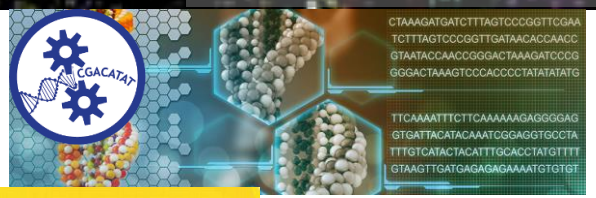

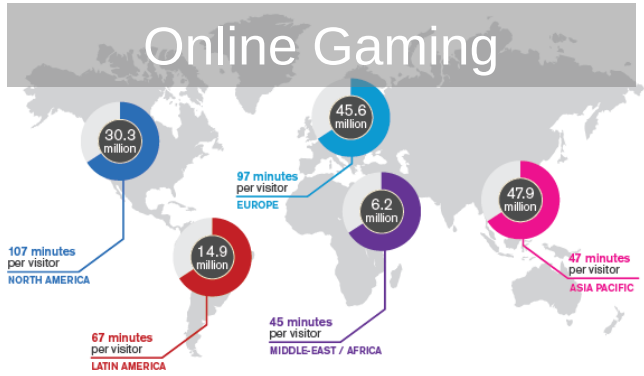

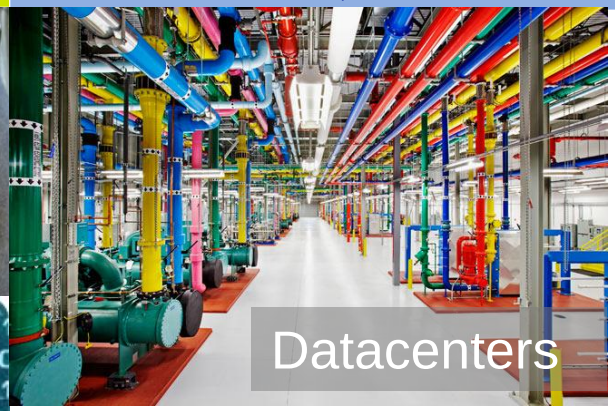

THE CURRENT TECHNOLOGY STACK: DATACENTER, SCHEDULER

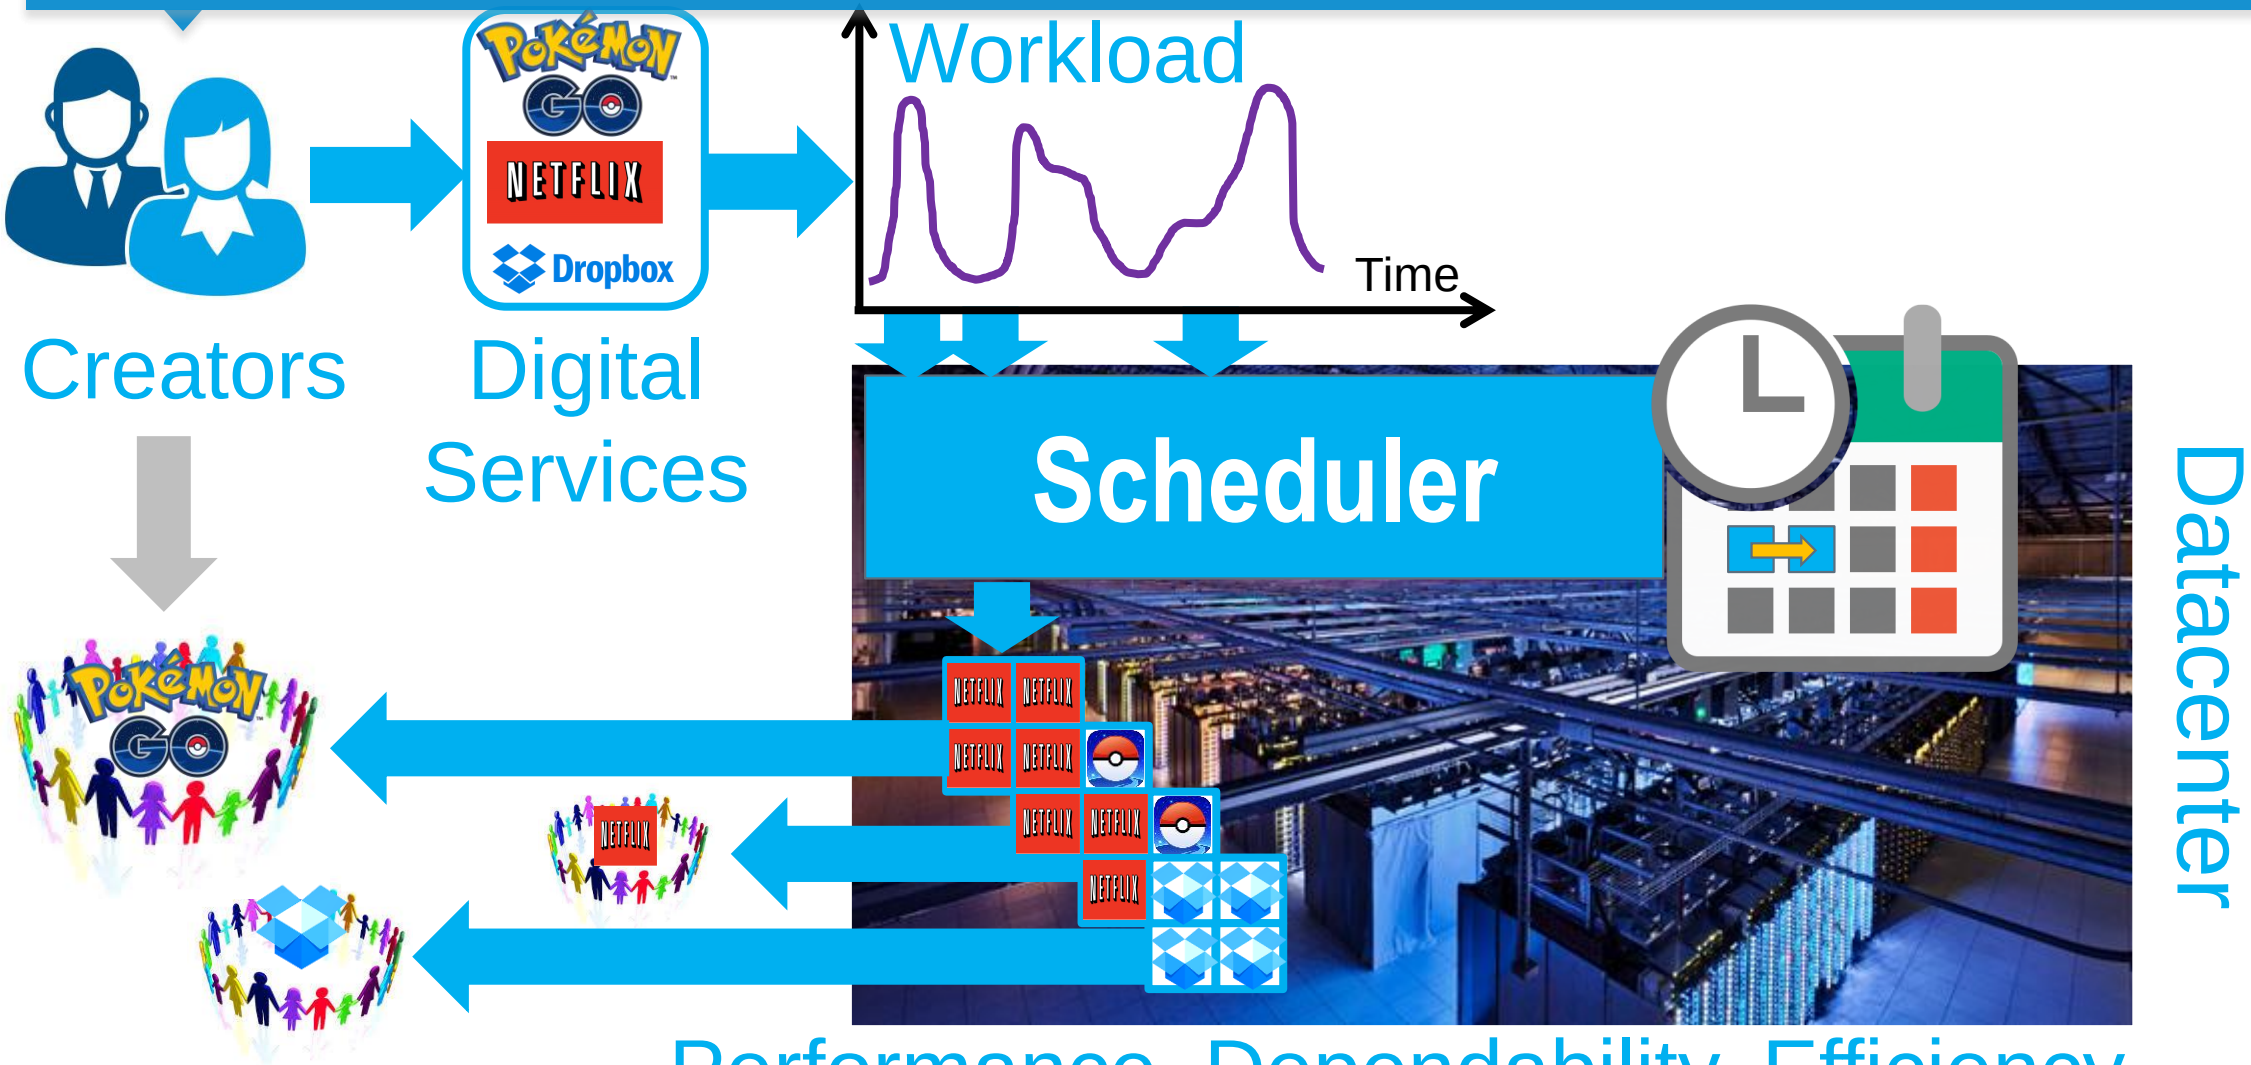

VISION: WHAT DOES OUR SOCIETY NEED?

PARTIALLY HAPPENING

"A world where individuals and human-centered organizations are augmented by an automated, sustainable layer of technology. At the core of this technology is ICT, and at the core of ICT are computer ecosystems, interoperating and performing as utilities and services, under human guidance and control."

- Clients:
  - > Fundamental right to ICT
  - > Understanding
- Computer professionals:
  - > Create
  - > Operate
  - > Experiment
  - > ...

- ICT:
  - > Utilities/services
  - > Automated
  - > Efficient
  - > Controlled
  - > Ecosystems
  - > Human-guided

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

IN THIS DIGITAL ECONOMY, FEW CAN BE SUCCESSFUL!

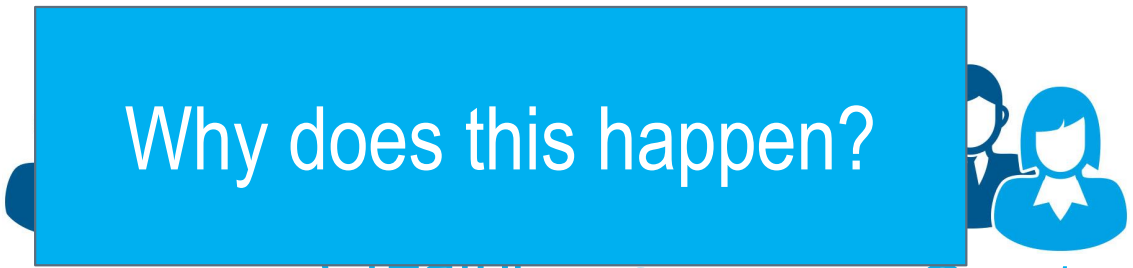

This Research

NETFLIX

Dropbox

CERN

"ICT is vital for SMEs", "SMEs are 60% GDP"

Sources: Eurostat'15,EC Digital Agenda,IDC'14

VU UNIVERSITY AMSTERDAM Faculty of Sciences

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

THE COMPLEXITY CHALLENGE

1. We Build and Test
Isolated Computer Systems, Yet
Everything Works in Ecosystems

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

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

PERFORMANCE, DEPENDABILITY, AND OTHER NON-FUNCTIONAL CHALLENGES

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

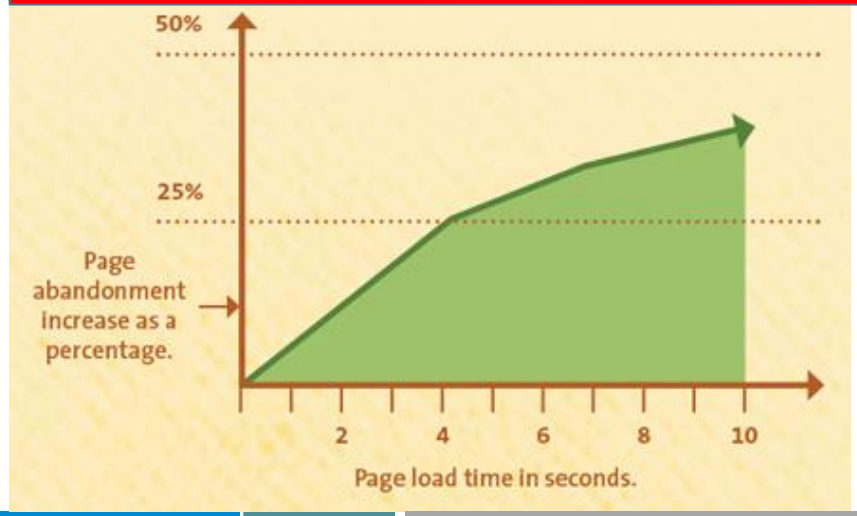

Google goes dark for 2 minutes, kills 40% of world's net traffic www.theregister.co.uk/2013/08/17/google_outage/

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 Also Educate Our Customers

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

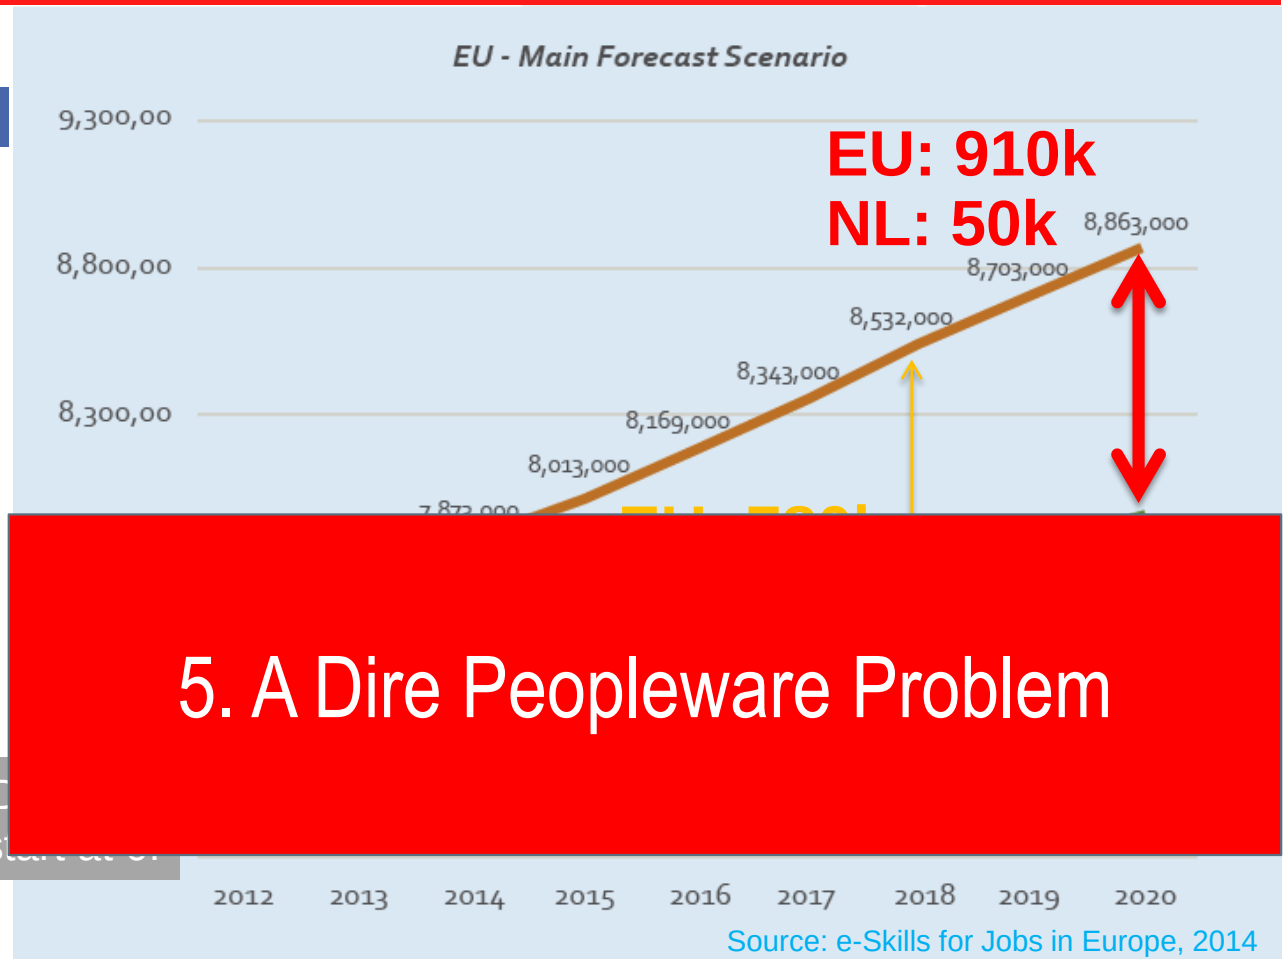

Source: e-Skills for Jobs in Europe, 2014

5. A Dire Peopleware Problem

VU UNIVERSITY AMSTERDAM Faculty of Sciences

THIS IS THE GOLDEN AGE OF DISTRIBUTED COMPUTER SYSTEMS

YET WE ARE IN A CRISIS – 5 CORE PROBLEMS

1. The Current Laws and Theories Are Built and Tested for Isolated Computer Systems

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 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 DISTRIBUTED SYSTEMS COURSE

Massivizing Computer Systems
The Science of Distributed Ecosystems

~40’

~10’ — Vision, Current Golden Age of Ecosystems, and Crisis

~5’ — Intermezzo: What Is Modern Science?

~20’ — Massivizing Computer Systems

1. Stakeholders
2. Meaningful discovery: science + engineering + design
3. Experimental method for discovery + Reproducibility/Validation
4. Codified and teachable Science
5. What’s left?

~2’ — Take-home message

BUT WHAT IS "A SCIENCE"? WHAT IS A MODERN SCIENCE?

Science starts from…

a new theory

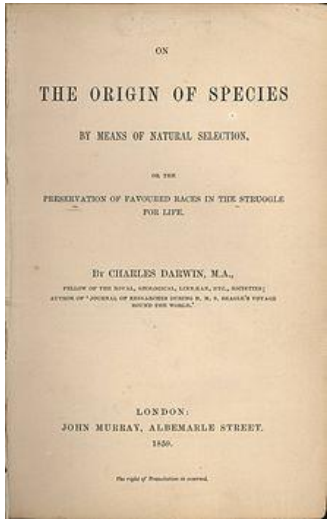

Darwin (1859)
The Origin of Species

a new way of thinking

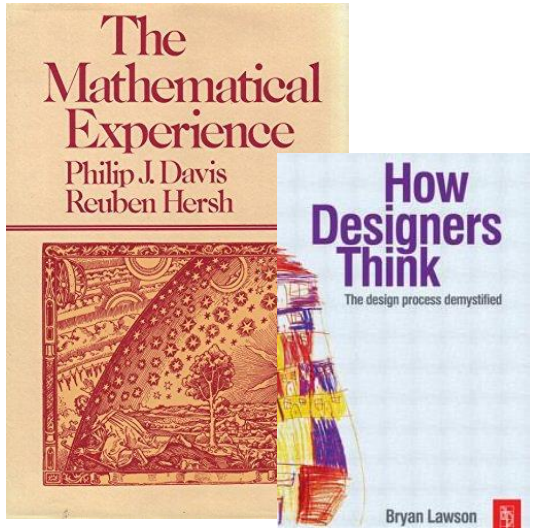

Davis and Hersh (1982)
The Mathematical Experience
Lawson (2005)
How Designers Think

several new principles

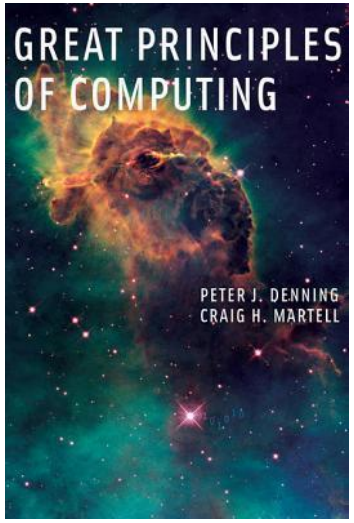

Denning and Martell (2015)
Great Principles of Computing

needs arise from practice

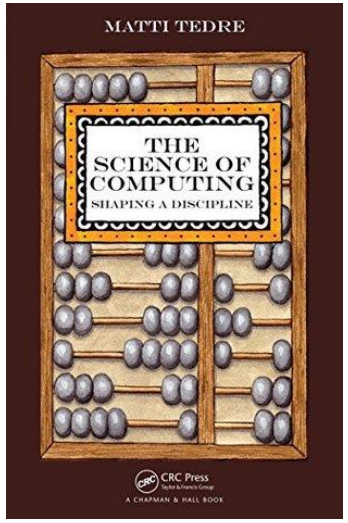

Tedre (2014)
The Science of Computing

theory-practice gap too large

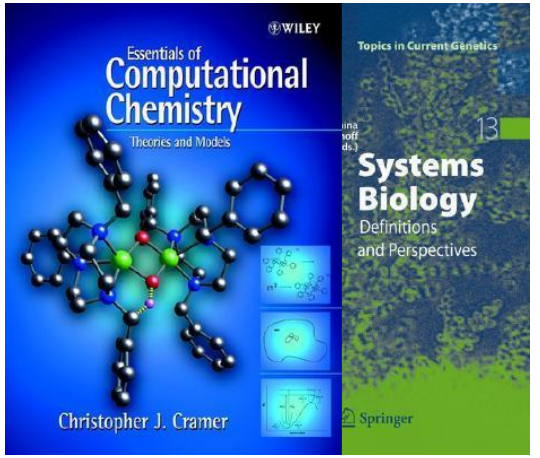

Cramer (2004) Essentials of Computational Chemistry
Alberghina and Westerhoff (2007)
Systems Biology

BUT WHAT IS "A SCIENCE"? WHAT IS A MODERN SCIENCE?

- Denning's criteria for a new field of science:
  - > Focus on pervasive phenomena
  - > By and for clear stakeholders
  - > Meaningful discovery
  - > Experimental methods of discovery
  - > Methods of validation
  - > Reproducibility of results
  - > Falsifiability of theoretical constructs
  - > Body of knowledge and skills that can be codified and taught

A SCIENCE OF ECOSYSTEMS

- My definition for a modern field of science:
  - > Fundamentals and Applications
  - > Methodology and Philosophy
  - > History and Envisioning
  - > Practice and Education
  - > Ethics

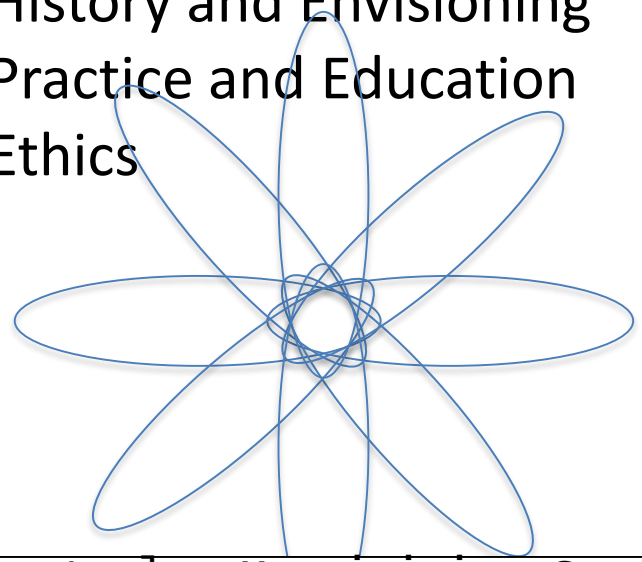

VU

Denning, The science in computer science, CACM 56, 2013.

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

Massivizing Computer Systems
The Science of Distributed Ecosystems

~40’

~10’ — Vision, Current Golden Age of Ecosystems, and Crisis

~5’ — Intermezzo: What Is Modern Science?

~20’ — Massivizing Computer Systems

1. Stakeholders
2. Meaningful discovery: science + engineering + design
3. Experimental method for discovery + Reproducibility/Validation
4. Codified and teachable Science
5. What’s left?

~2’ — Take-home message

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]

~40’

Massivizing Computer Systems
The Science of Distributed Ecosystems

~10’ — Vision, Current Golden Age of Ecosystems, and Crisis

~5’ — Intermezzo: What Is Modern Science?

~20’ — Massivizing Computer Systems

1. Stakeholders
2. Meaningful discovery: science + engineering + design
3. Experimental method for discovery + Reproducibility/Validation
4. Codified and teachable Science
5. What’s left?

~2’ — Take-home message

STAKEHOLDERS ALSO INCLUDE: ATLARGE RESEARCH, OUR TEAM

http://atlarge.science/people.html

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

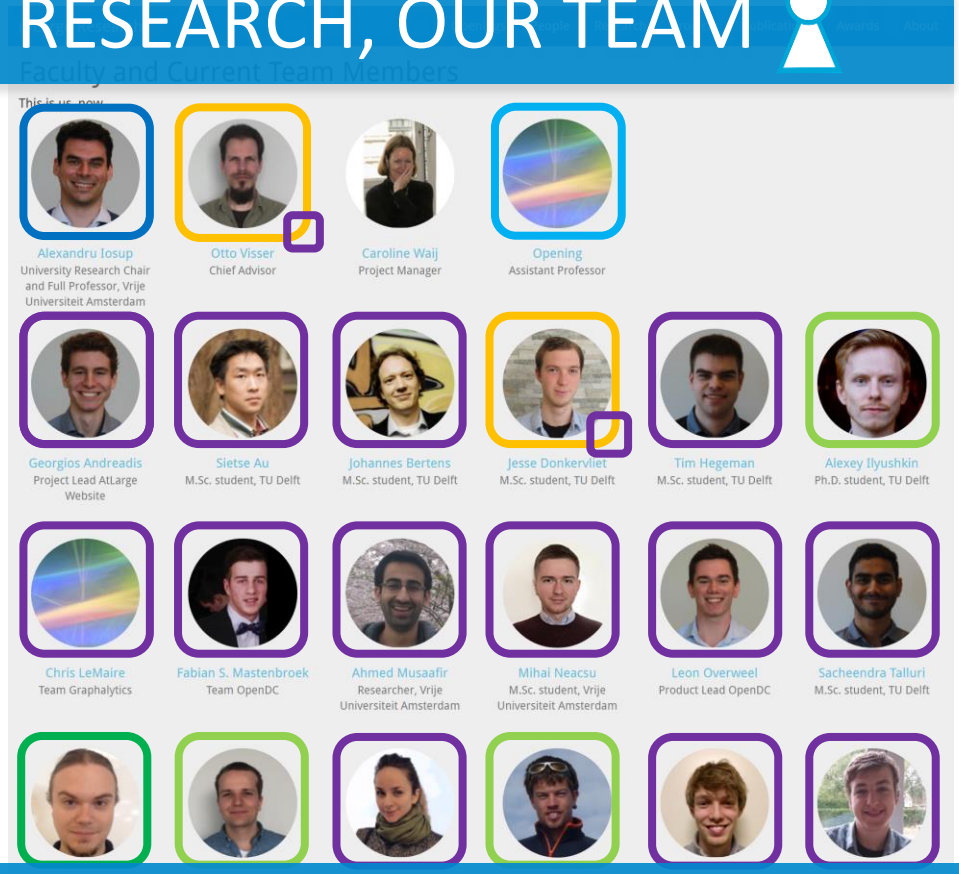

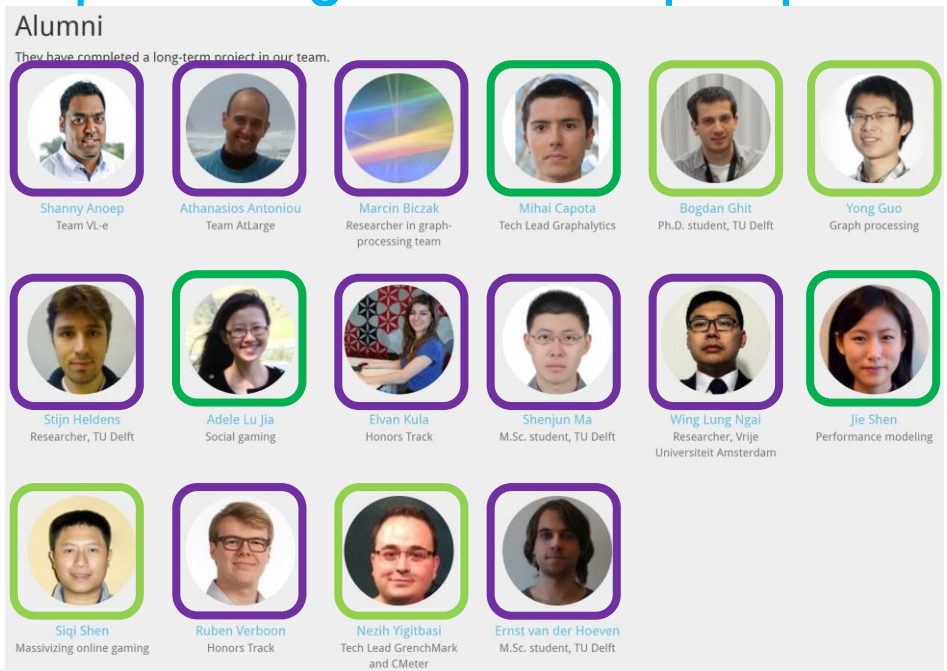

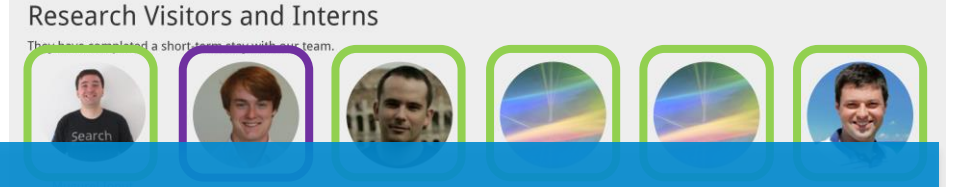

WE ARE A DIVERSE GROUP, OF DIFFERENT RACES AND ETHNICITIES, GENDERS AND SEXUAL PREFERENCES, VIEWS OF CULTURE , POLITICS, AND RELIGION. YOU ARE WELCOME TO JOIN!

~40’

Massivizing Computer Systems
The Science of Distributed Ecosystems

~10’ — Vision, Current Golden Age of Ecosystems, and Crisis

~5’ — Intermezzo: What Is Modern Science?

~20’ — Massivizing Computer Systems

1. Stakeholders
2. Meaningful discovery: science + engineering + design
3. Experimental method for discovery + Reproducibility/Validation
4. Codified and teachable Science
5. What’s left?

~2’ — Take-home message

MASSIVIZING COMPUTER SYSTEMS: MEANINGFUL DISCOVERY

Scheduling
Workflows
Domain-Specific/Agnostic
Portfolio, Auto-scaling*

Dependability
Performance & Failure Analysis*
Space-/Time-Correlation
Availability-On-Demand

New World
Workload Modeling
Serverless
Reference Architectures

Ecosystem Navigator
Performance Variability
Grid*, Cloud, Big Data
Benchmarking*
Longitudinal Studies

Scalability/Elasticity
Delegated Matchmaking*
BTWorld*, POGGI*, AoS
Auto-Scalers*
Heterogeneous Systems

Socially Aware
Collaborative Downloads*
Groups in Online Gaming
Toxicity Detection*
Interaction Graphs

Education
Social Gamification*

Software Artifacts
Graphalytics, OpenDC

Data Artifacts
Distributed Systems Memex*

Fundamental Problems/Research Lines
Our Contribution So Far

* Award-level work
Competitive personal grants

MASSIVIZING COMPUTER SYSTEMS: MEANINGFUL DISCOVERY

Scheduling
Workflows

Dependability
Performance & Failure Analysis*

New World
Workload Modeling

Many thanks to all collaborators* for making this research possible:

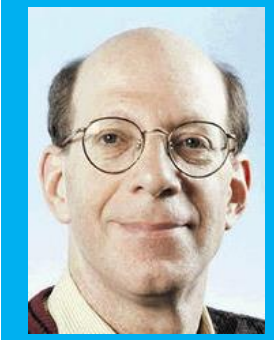
Andy


Dror


Sam


Miron


Dick


Fernando


Radu


Ion


Ana

* Including all collaborators, I have published with over 180 people. Of them, about 40 were part of the team in the Netherlands.

Education
Social Gamification*

Software Artifacts
Graphalytics, OpenDC

Data Artifacts
Distributed Systems Memex*

Fundamental Problems/Research Lines

Our Contribution So Far

* Award-level work

Competitive personal grants

MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE

GALILEO GALILEI, 1608-9, 3-8X TELESCOPE

MERELY AN INSTRUMENT?

FUNDAMENTAL SCIENCE?

Garney. The Inquisition's Semicolon: Punctuation, Translation, and Science in the 1616 Condemnation of the Copernican System, ArXiv document 1402.6168. [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

400 YEARS LATER, SYSTEMATIC SURVEY

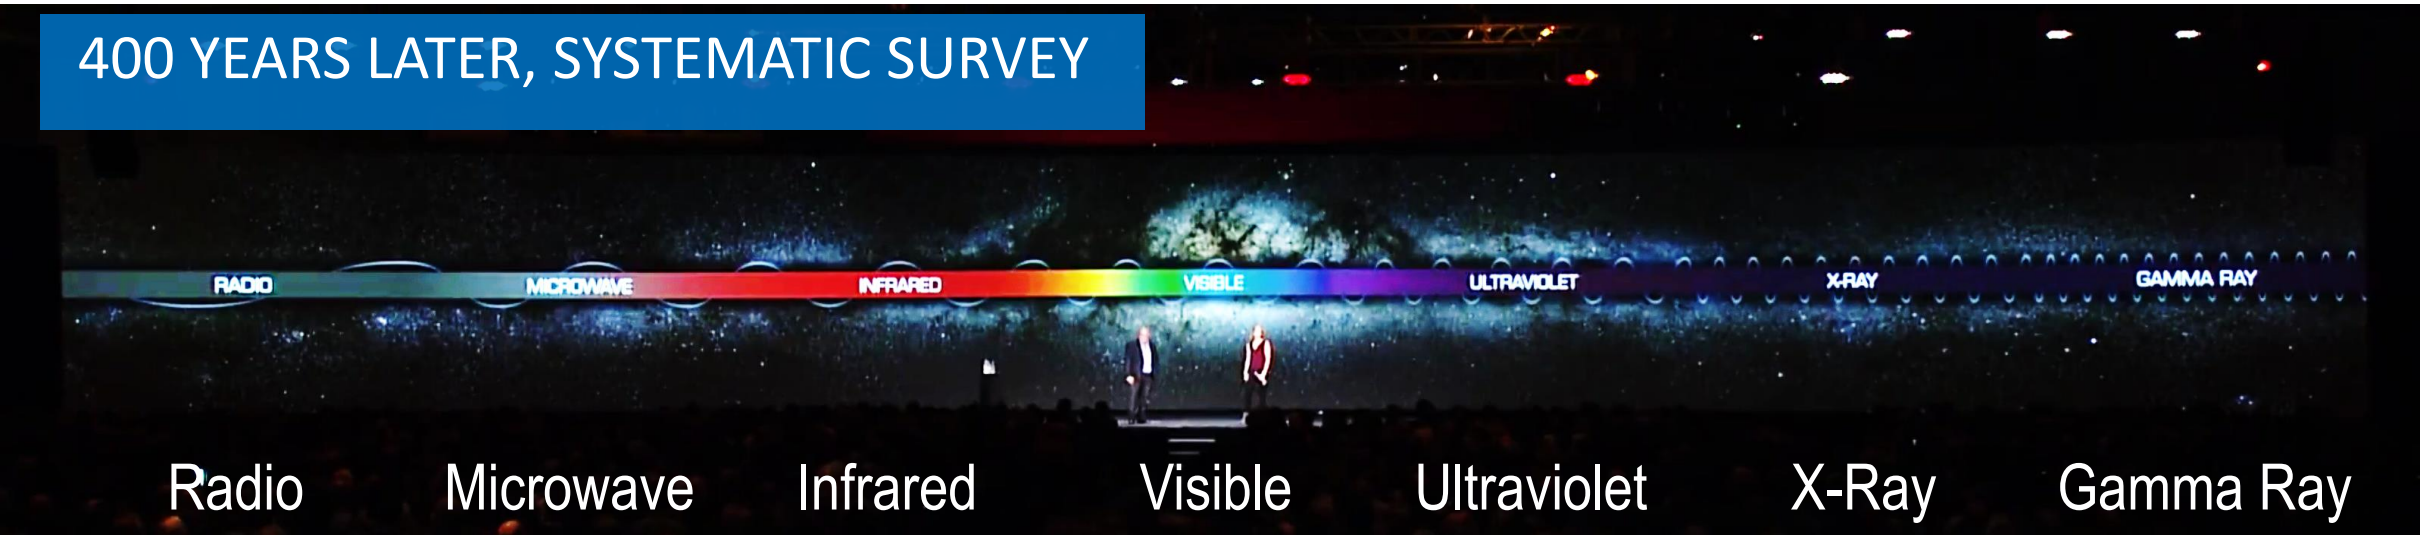

MEANINGFUL DISCOVERY

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

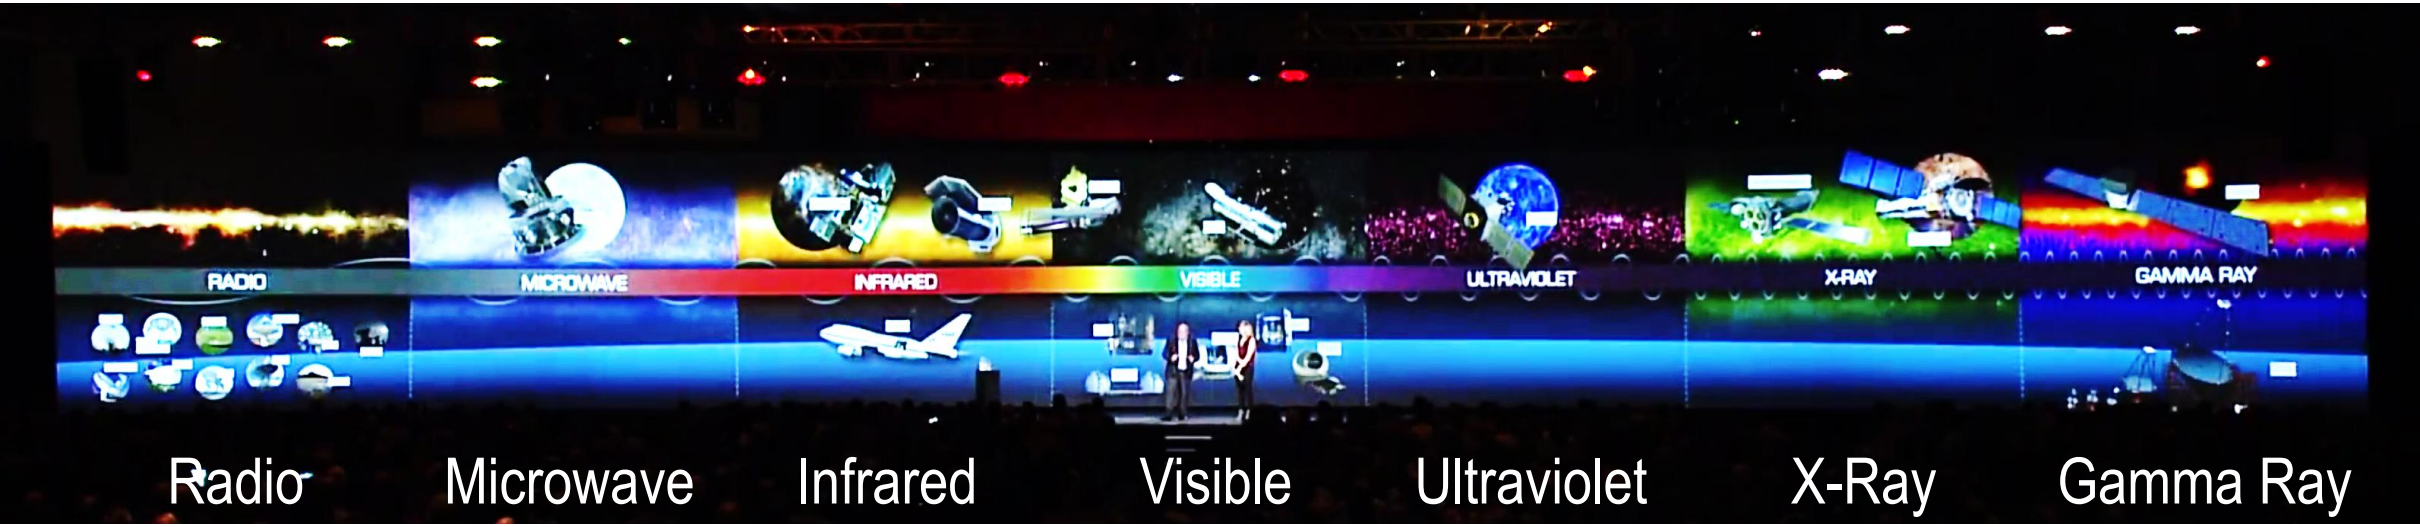

Radio	Microwave	Infrared	Visible	Ultraviolet	X-Ray	Gamma Ray
Cloud, Grid, Edge, Fog, etc.	Big Data	Sci.&Eng. Apps	Consumer Apps	Enterprise Apps	Systems, Ecosystems	Performance, Security, etc.

MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

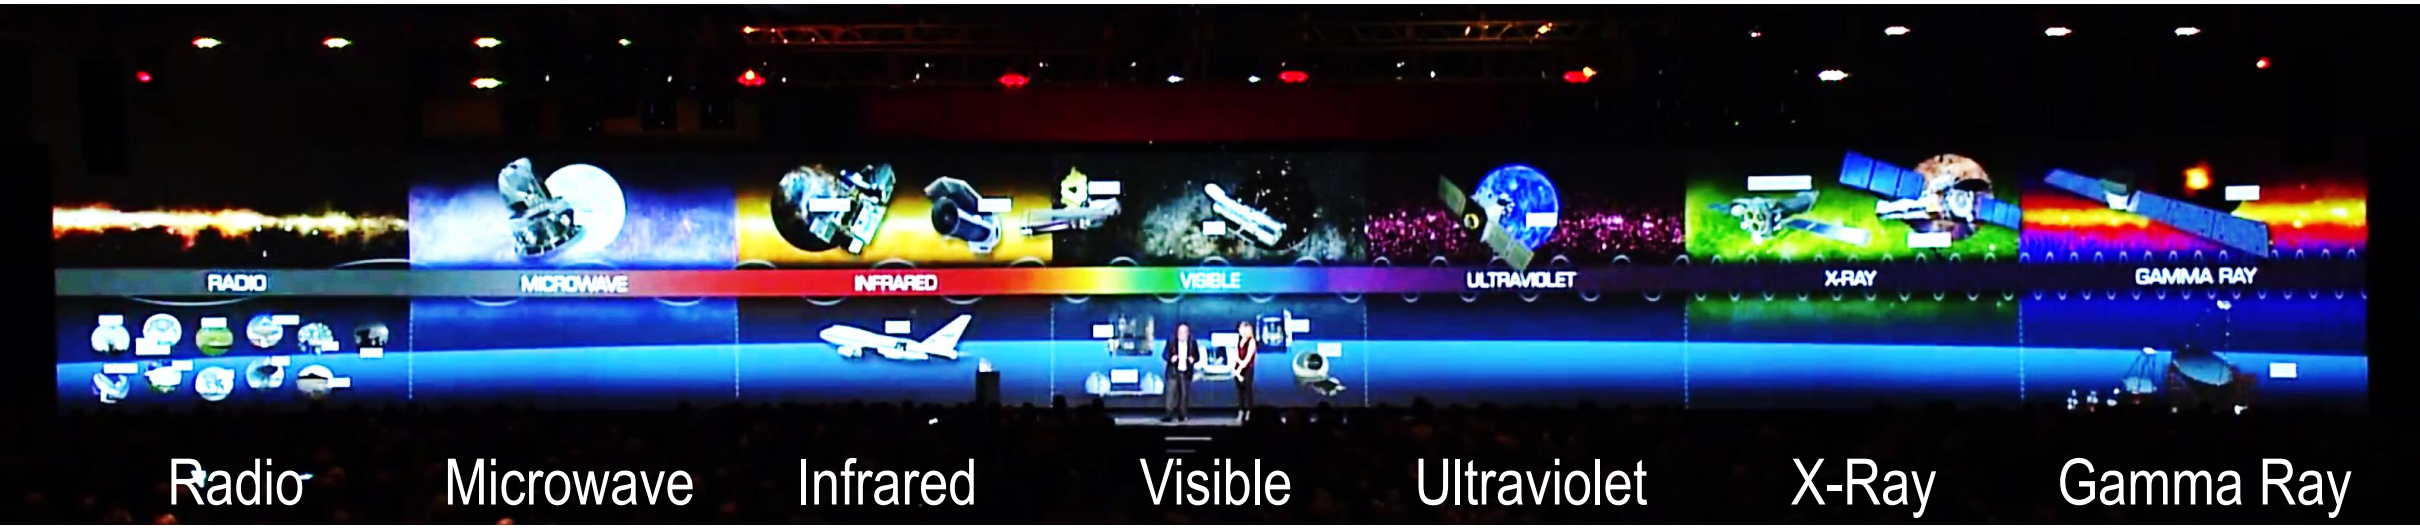

Radio	Microwave	Infrared	Visible	Ultraviolet	X-Ray	Gamma Ray
Cloud, Grid, Edge, Fog, etc.	One aspect: BigData, P2P	Sci.&Eng. Apps+Sys.	Consumer Apps+Sys.	Enterprise Sys.	Systems, Ecosystems	Performance, Availability, etc.
						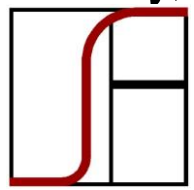
[Iosup et al. FGCS'08]	[Zhang et al. CoNext'10]	[Iosup et al. IEEE IC'11]	[Guo et al. NETGAMES'12]	[Shen et al. CCGRID'15]	[Ghit et al. CCGRID'14]	[Iosup et al. CCGRID'10]

MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

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

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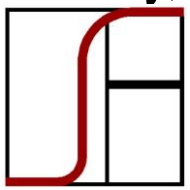
[Iosup et al. FGCS'08]	[Zhang et al. CoNext'10]	[Iosup et al. IEEE IC'11]	[Guo et al. NETGAMES'12]	[Shen et al. CCGRID'15]	[Ghit et al. CCGRID'14]	[Iosup et al. CCGRID'10]

MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

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

…

BOTS, NOT PARALLEL JOBS

Cloud, Grid, Edge, Fog, etc.

[Iosup et al. FGCS'08]

…

GROUPS NOT RARE, DOMINANT

One aspect: BigData, P2P

[Zhang et al. CoNext'10]

…

COMMUNITY FORMATION

Sci.&Eng. Apps+Sys.

[Iosup et al. IEEE IC'11]

…

SYSTEMIC VARIABILITY

Consumer Apps+Sys.

[Guo et al. NETGAMES'12]

…

CORRELATED, NOT IID FAILURES

Enterprise Sys.

[Shen et al. CCGRID'15]

Systems, Ecosystems

[Ghit et al. CCGRID'14]

Performance, Availability, etc.

[Iosup et al. CCGRID'10]

MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

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

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.
		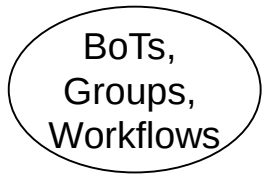				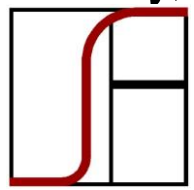
[Iosup et al. FGCS'08]	[Zhang et al. CoNext'10]	[Iosup et al. IEEE IC'11]	[Guo et al. NETGAMES'12]	[Shen et al. CCGRID'15]	[Ghit et al. CCGRID'14]	[Iosup et al. CCGRID'10]

MEANINGFUL DISCOVERY

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

FUNDING: > 3B USD

- The Human Genome Project:
  - > Physical map covering >90% human genome
  - > Sequence data made available open-access
- Big Science:
  - > Took >10 years to complete
  - > Led by US, work by 20 groups in CN, DE, FR, JP, UK, US
- Big impact:
  - > Decrease cost of sequencing
  - > Facilitate biomedical research

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

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

REMEMBER THE COMPLEXITY CHALLENGE?

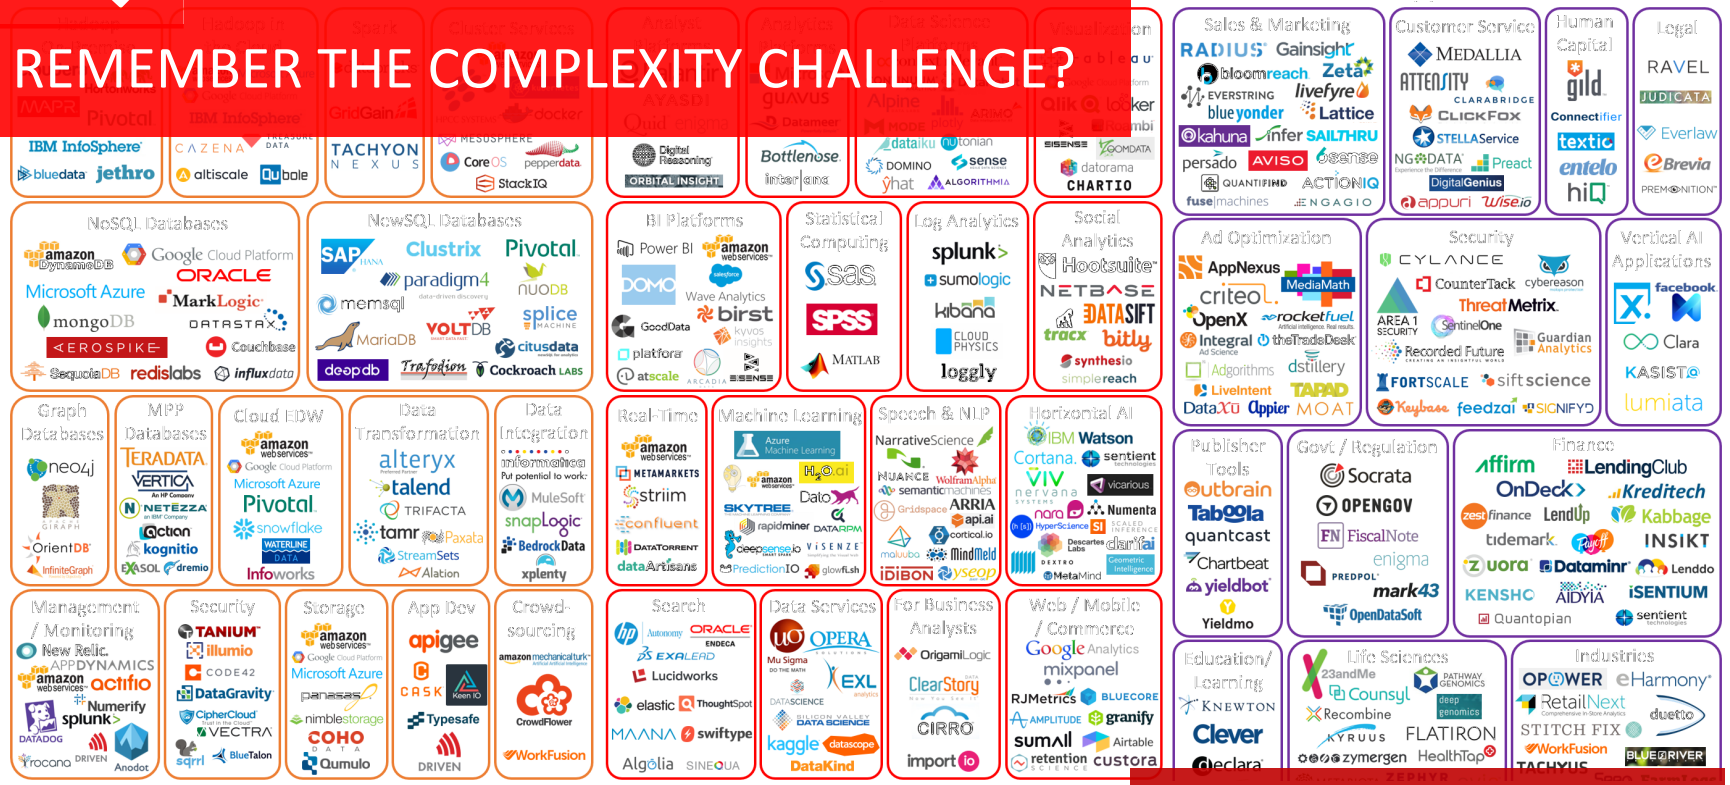

<<1% OF BIG DATA
BY MATT TURK (2017)

MEANINGFUL DISCOVERY

THE COMPLEXITY CHALLENGE

Focus on Applications, 5 Core Layers:

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

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS

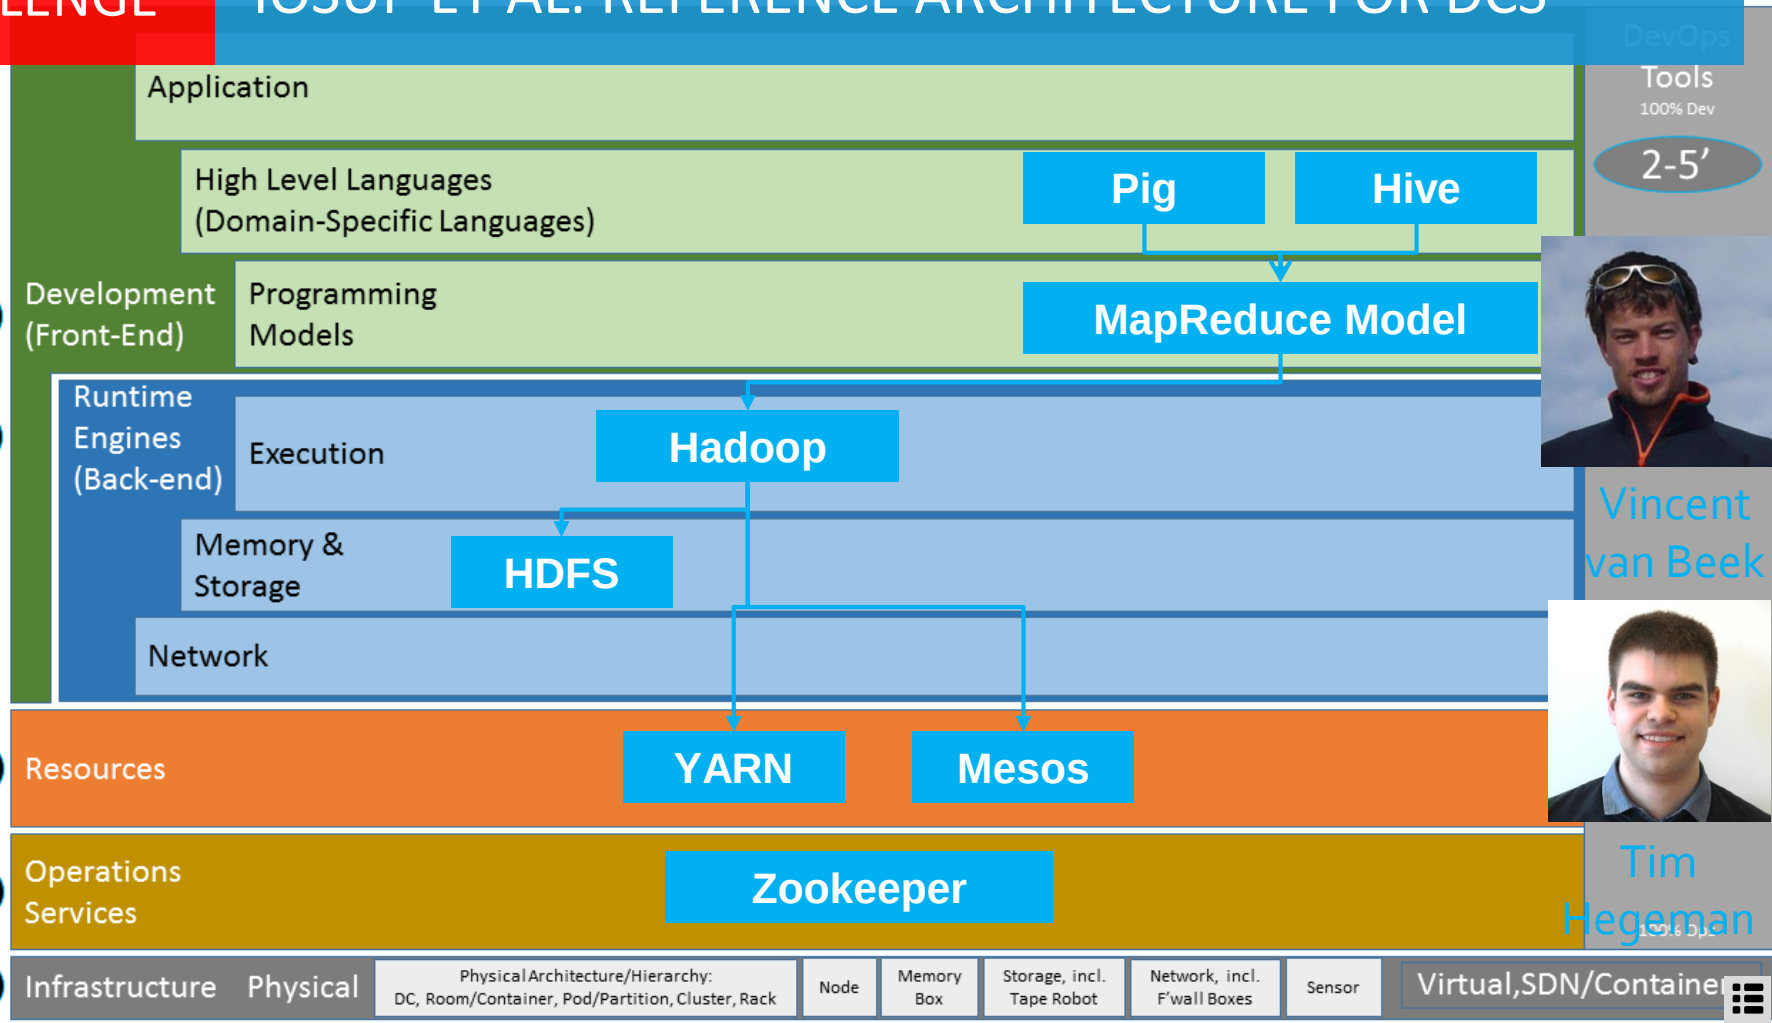

Vincent van Beek

Tim Hegeman

VU

[Iosup et al. ICDCS'18]

MEANINGFUL DISCOVERY

THE COMPLEXITY CHALLENGE

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS

Georgios Andreadis

ANDREADIS ET AL. REFERENCE ARCHITECTURE FOR SCHEDULERS IN DCS

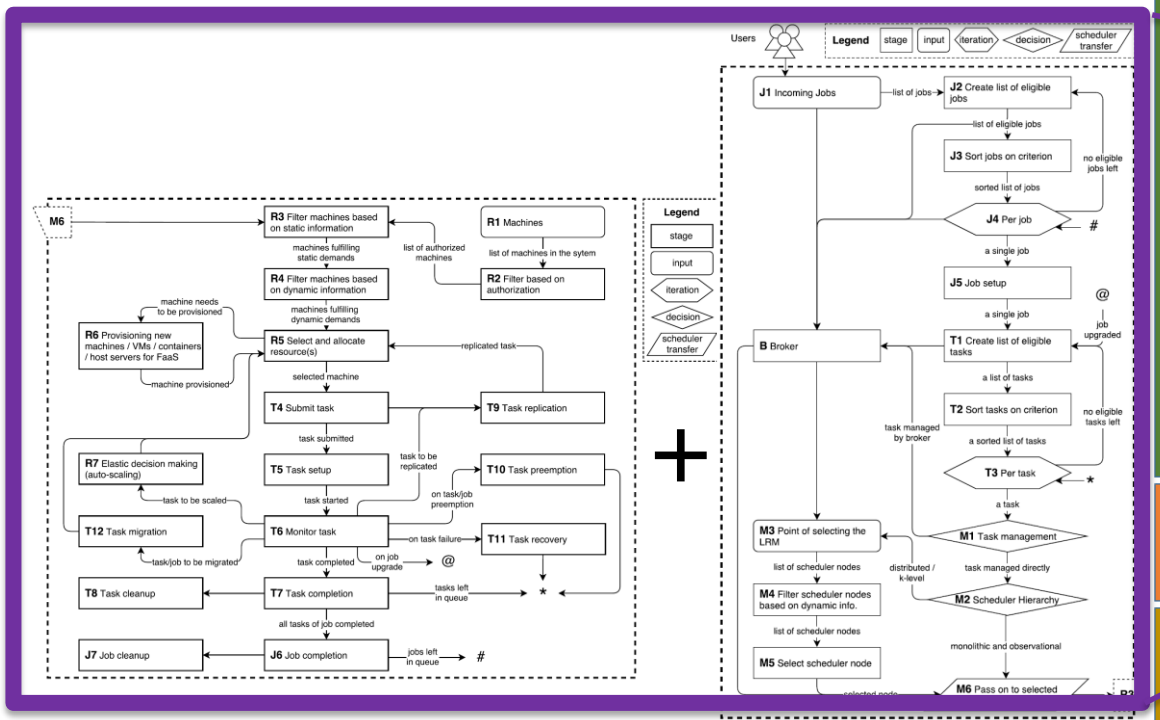

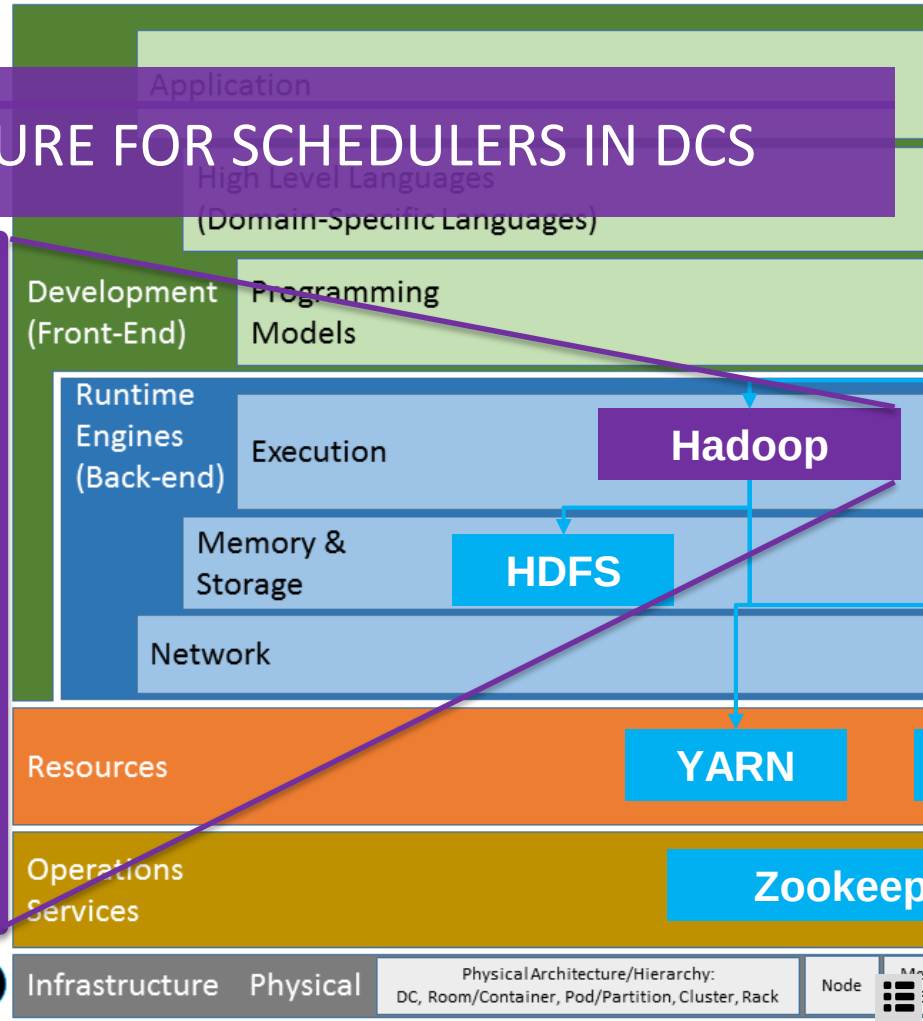

[Andreadis et al. SC'18]

MEANINGFUL DISCOVERY

science + engineering + design

[Iosup et al. ICDCS'18]

MEANINGFUL DISCOVERY

NO SYSTEMATIC PROCESS FOR COMPUTER SYSTEMS

SO I'LL USE EXAMPLES

science + engineering + design

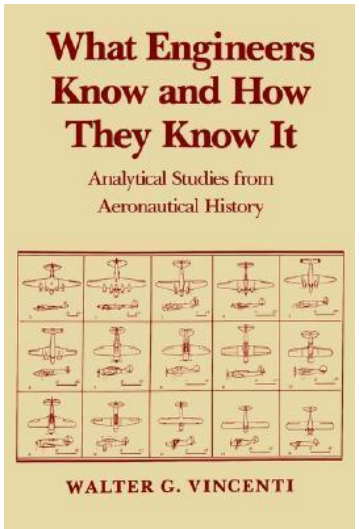

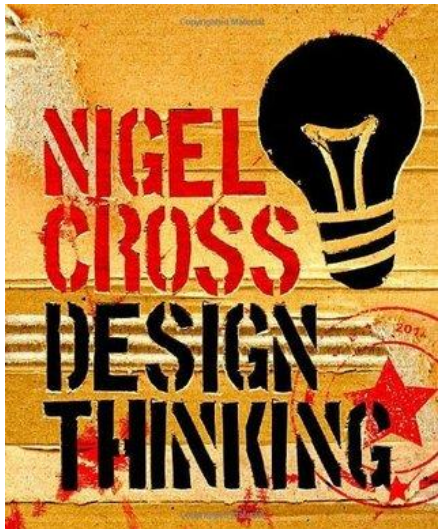

THE COMPUTER SYSTEMS TRIPLET

[Iosup et al. ICDCS'18]

MEANINGFUL DISCOVERY

ENGINEERING LDBC GRAPHALYTICS: THE NEED FOR SPEED … & GRAPHS!

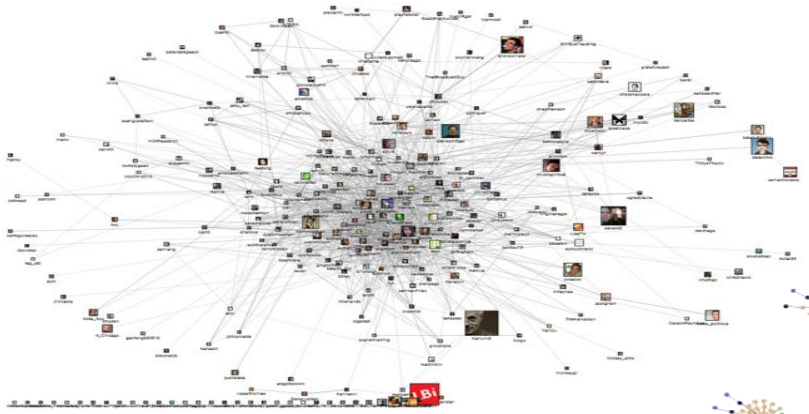

Social network

~1 billion vertices
~100 billion connections

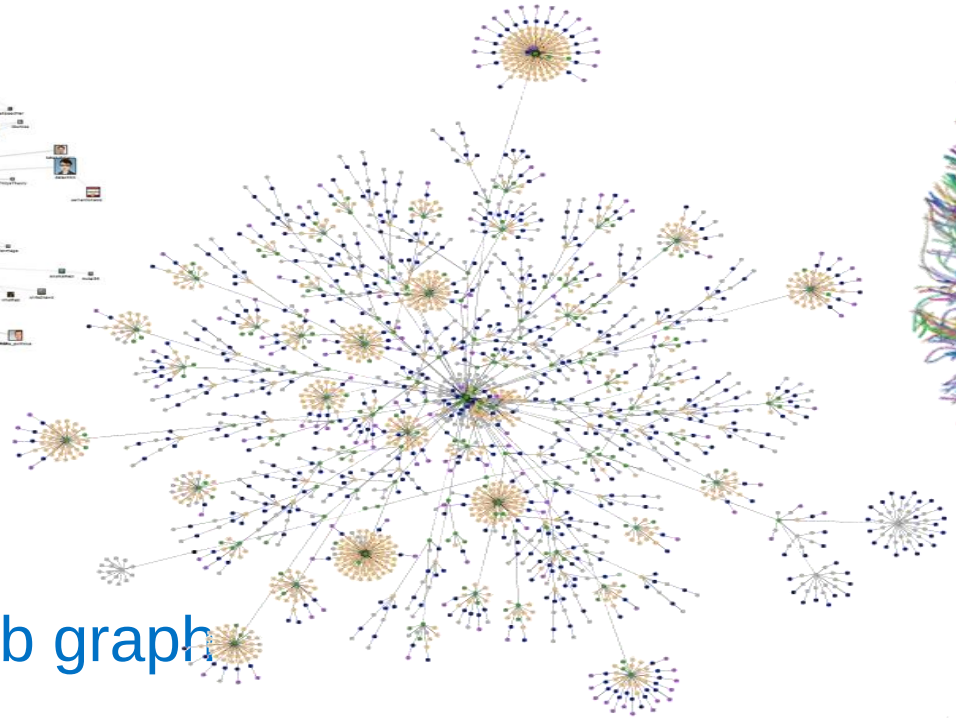

Web graph

~50 billion pages
~1 trillion hyperlinks

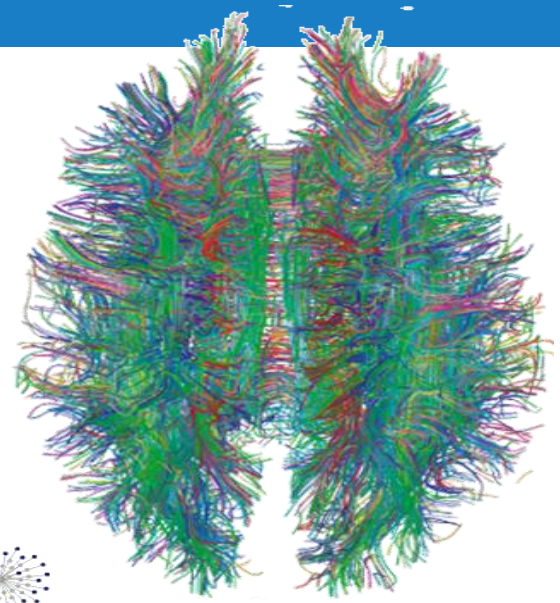

Brain network

~100 billion neurons
~100 trillion connections

Sources: Smith, CHI'10; Blog webpage; Gigandet et al., PLoS ONE 3(12)]

MEANINGFUL DISCOVERY

ENGINEERING LDBC GRAPHALYTICS: THE SYSTEMS LANDSCAPE

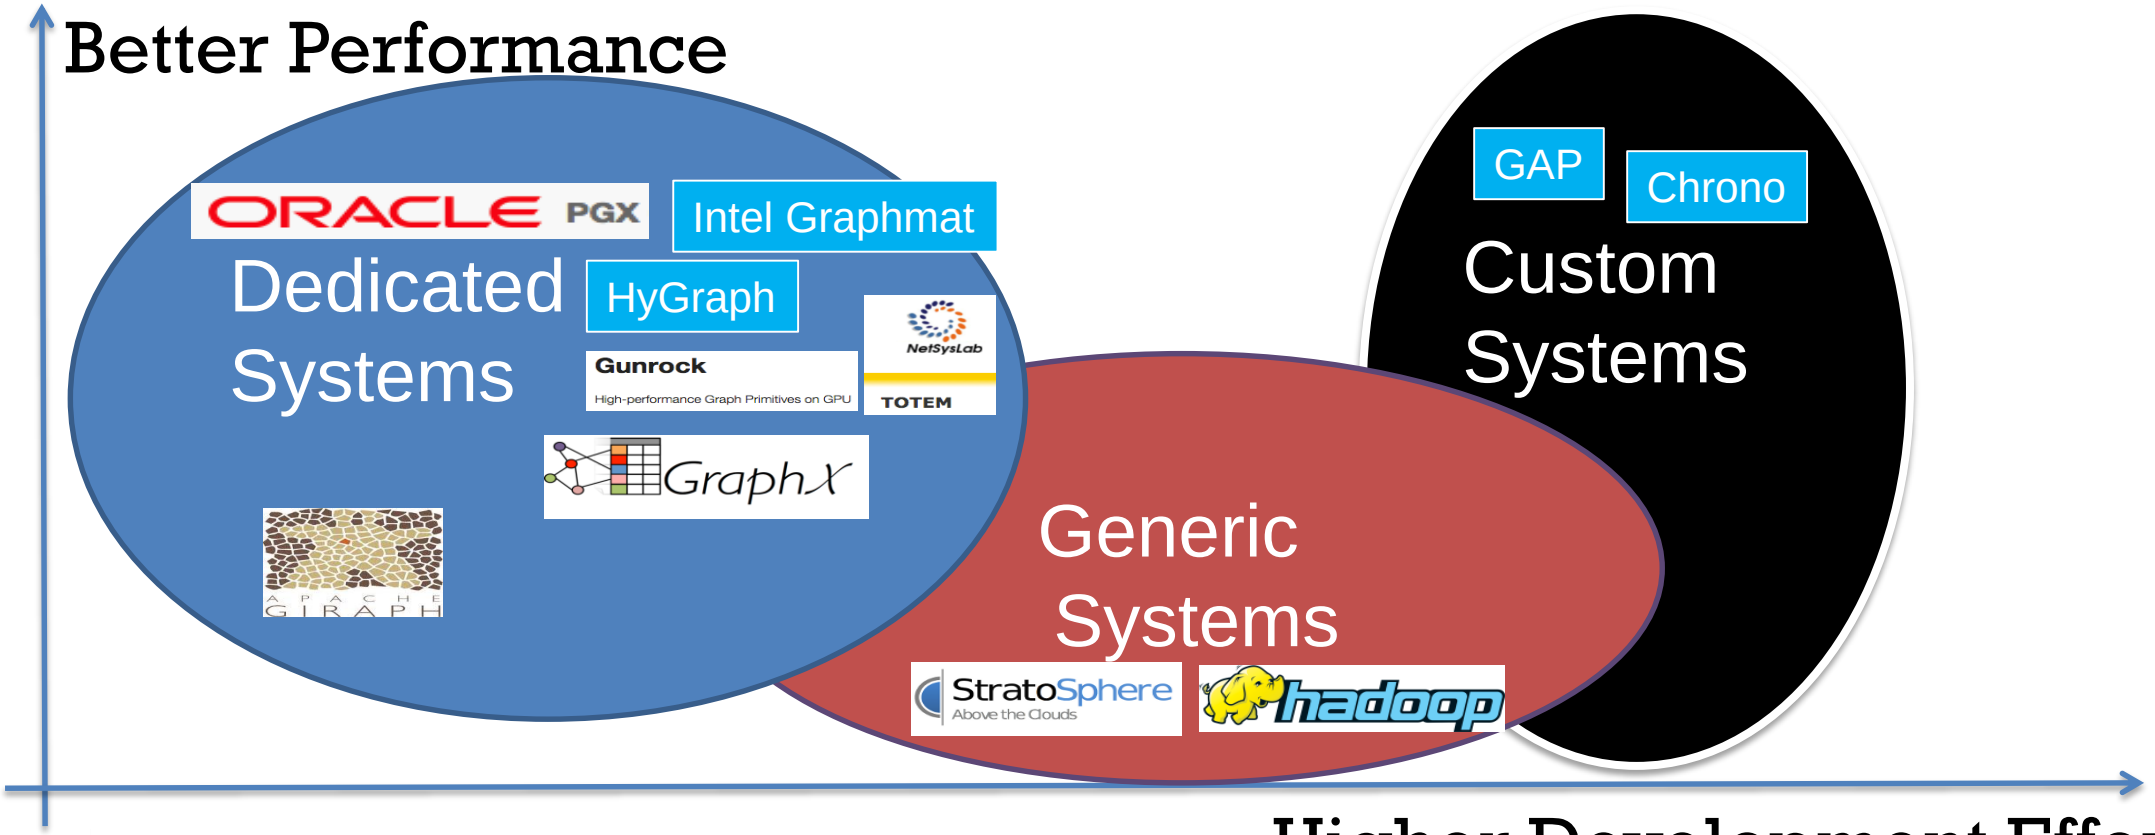

MEANINGFUL DISCOVERY

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, derived from practice
  - > Renewal process to keep the workload relevant
  - > Enables comparison of many platforms, community-driven and industrial
  - > Global Competition

Wing Lung Ngai

Tim Hegeman

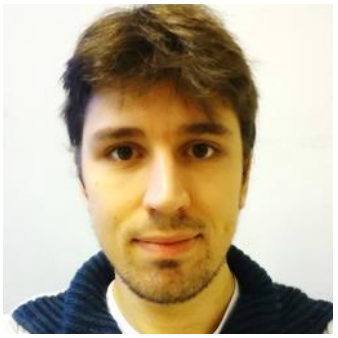

Stijn Heldens

Alex Uță

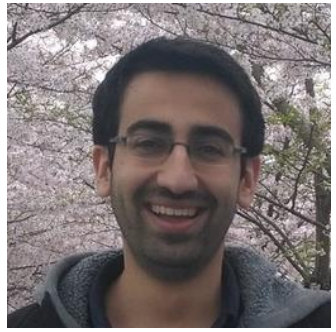

Ahmed Musaafir

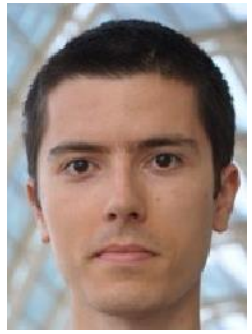

Mihai Capotă

[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

ENGINEERING LDBC GRAPHALYTICS: MODELING LEADS TO PERFORMANCE ANALYSIS

LDBC The graph & RDF benchmark reference

- Graphalytics Grade10:
  - > Automated bottleneck detection
  - > Automated identification of performance issues

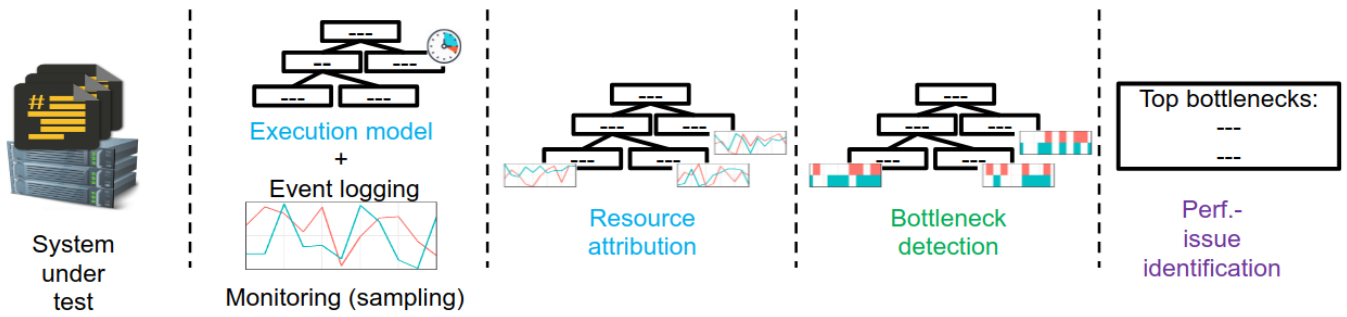

Multi-stage process, works in ecosystem

Tim Hegeman

MEANINGFUL DISCOVERY

ENGINEERING LDBC GRAPHALYTICS: MODELING LEADS TO PERFORMANCE ANALYSIS

LDBC The graph & RDF benchmark reference

- Graphalytics Grade10:
  - > Automated bottleneck detection
  - > Automated identification of performance issues

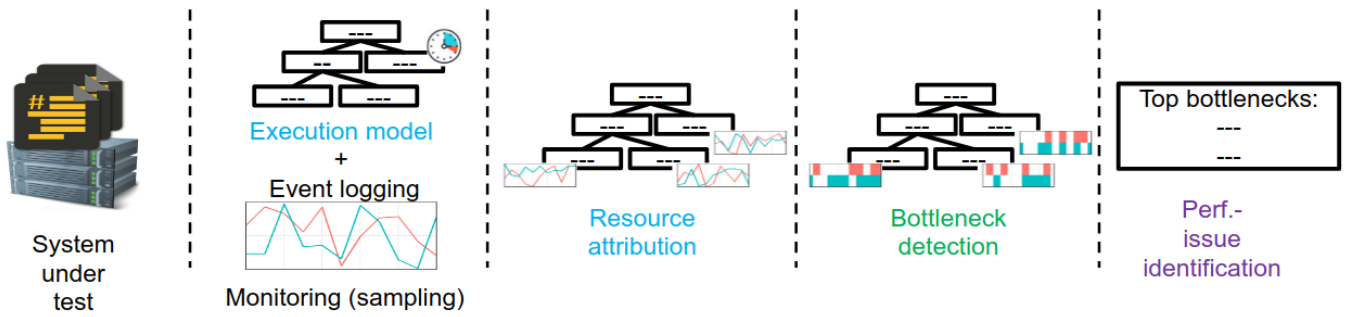

Multi-stage process, works in ecosystem

APACHE GIRAPH

- Without Grade10:

No bottleneck at all

- With Grade10:

Always bottleneck
Cause:
+ Message queue full
+ Garbage collector
+ CPU
+ Others

MEANINGFUL DISCOVERY

DESIGNING SERVERLESS ARCHITECTURES

Abstraction: Serverless
Design: FaaS systems

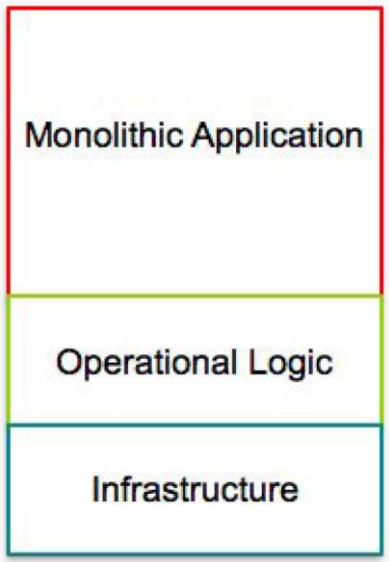

Difficult to Scale
Infrequent, Inflexible
Complex deployment
Tightly coupled stack

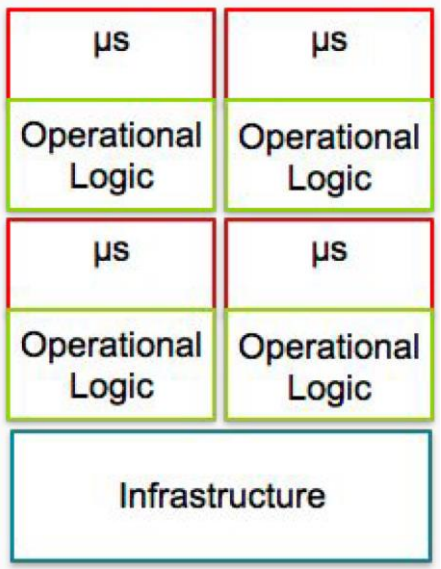

- Scalable
- Frequent, Flexible
- Complexity: from application logic to operational logic

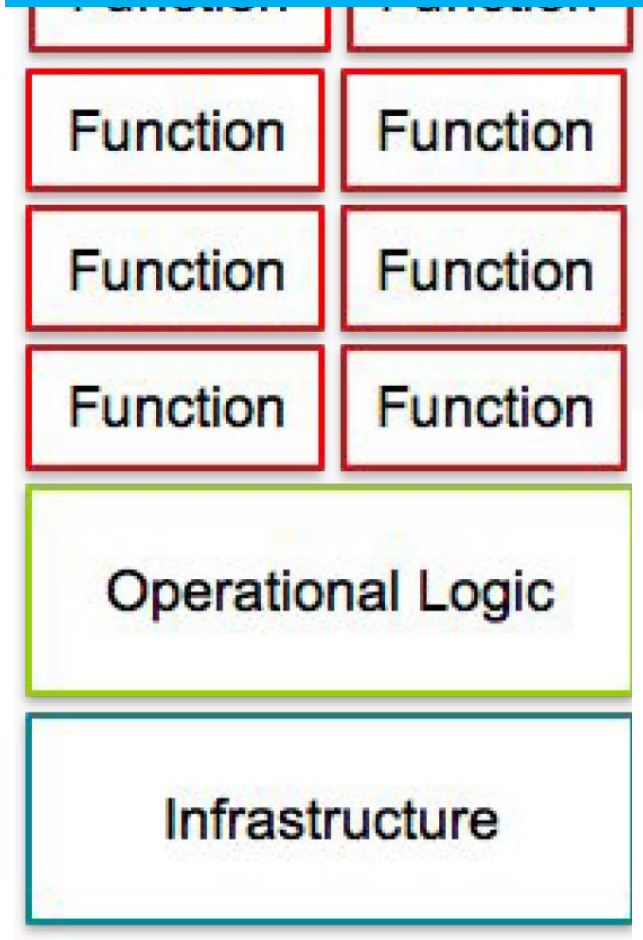

Erwin van Eyk

Lucian Toader

- Scalable
- Frequent, Flexible
- Explicit separation of Business Logic vs. Operational Logic.
- Minimal layer coupling, unit of deployment

MEANINGFUL DISCOVERY

DESIGNING SERVERLESS ARCHITECTURES

Erwin
van Eyk

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

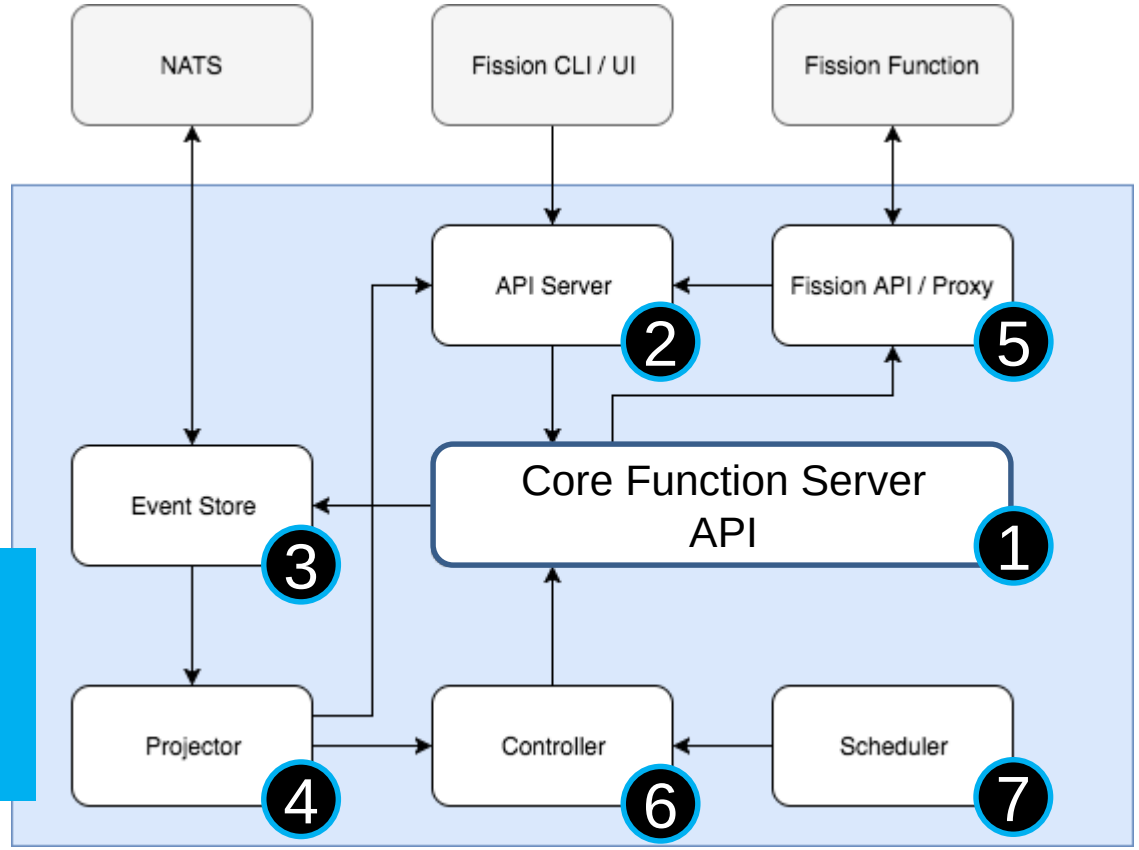

https://github.com/fission/fission-workflows/blob/master/Docs/architecture.md

Massivizing Computer Systems
The Science of Distributed Ecosystems

~40’

~10’ — Vision, Current Golden Age of Ecosystems, and Crisis

~5’ — Intermezzo: What Is Modern Science?

~20’ — Massivizing Computer Systems

1. Stakeholders
2. Meaningful discovery: science + engineering + design
3. Experimental method for discovery + Reproducibility/Validation
4. Codified and teachable Science
5. What’s left?

~2’ — Take-home message

EXPERIMENTAL METHODS OF DISCOVERY

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

Our Prototypes (*in phyisico/in vitro*)

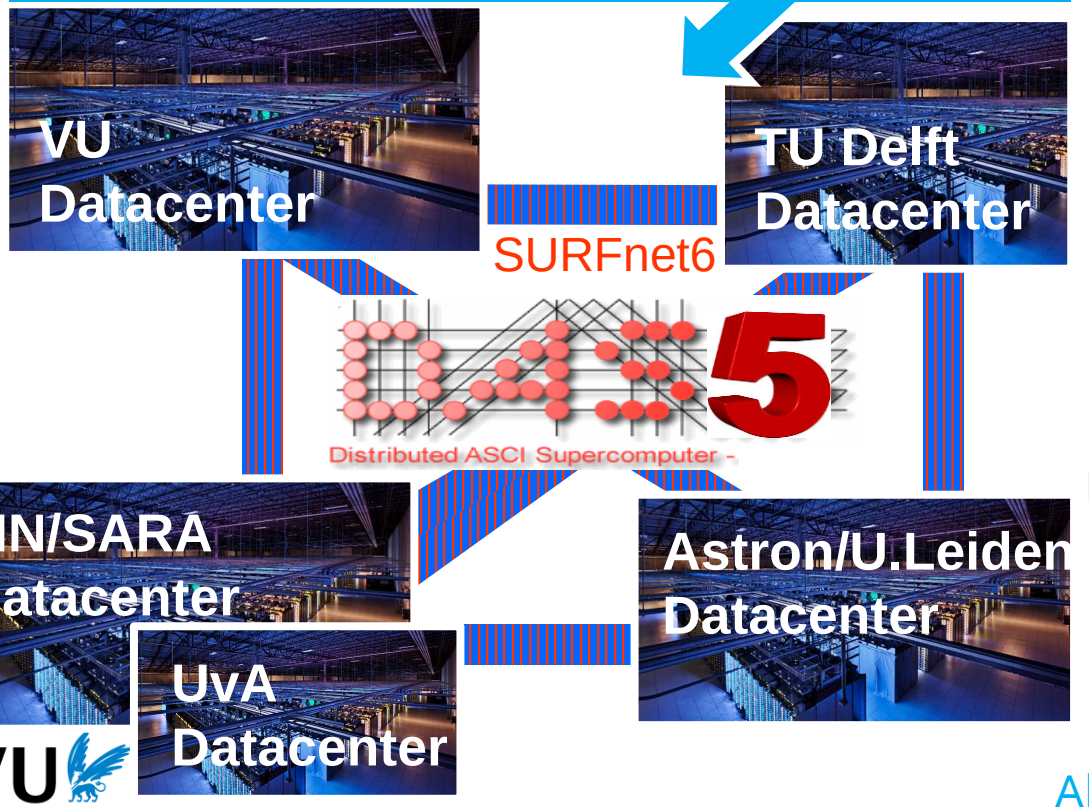

Laurens Versluis

Georgios Andreadis

Fabian Mastebroek

Lacheendra

Maria Voinea

Alexey Ilyushkin

\+

amazon web services™

We also use clouds

\+ OpenDC

And simulators (*in silico*)

REPRODUCIBILITY AND VALIDATION OF DISCOVERY

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

METHODOLOGY	SHARED PRINCIPLES, METHODS, ETC. … BUT WHERE*?!
OPEN SCIENCE	FREE OPEN-SOURCE SOFTWARE AND OPEN-ACCESS DATA … BUT WHO PAYS/ WHAT INCENTIVES?!
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…

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

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

~40’

Massivizing Computer Systems
The Science of Distributed Ecosystems

~10’ — Vision, Current Golden Age of Ecosystems, and Crisis

~5’ — Intermezzo: What Is Modern Science?

~20’ — Massivizing Computer Systems

1. Stakeholders
2. Meaningful discovery: science + engineering + design
3. Experimental method for discovery + Reproducibility/Validation
4. Codified and teachable Science
5. What’s left?

~2’ — Take-home message

21ST CENTURY NEEDS FOR MASSIVIZING SYSTEMS

AN EDUCATION FRAMEWORK WITH FOUR DIMENSIONS

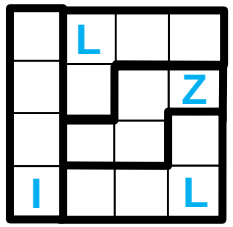

The Other Path

21ST CENTURY NEEDS FOR MASSIVIZING SYSTEMS

AN EDUCATION FRAMEWORK WITH FOUR DIMENSIONS

1. Curriculum: what to teach?
2. Didactics: how to teach?
3. Technology: how to address higher and more diverse expectations?
4. Management: how to manage, incentivize also the educators?

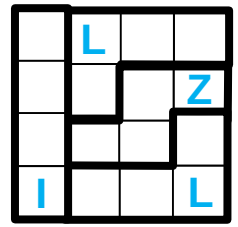

The Other Path

21ST CENTURY NEEDS FOR MASSIVIZING SYSTEMS

DIDACTICS: SOCIAL GAMIFICATION AS INTUITION AND METHOD

Gamification = The use of thinking and techniques designed for gaming in non-gaming settings, e.g., in education.

http://goo.gl/V97zSW

What is the intuition behind gamification?

How can gamification be used?

http://goo.gl/ILSNeb

21ST CENTURY NEEDS FOR MASSIVIZING SYSTEMS

DIDACTICS: SOCIAL GAMIFICATION WORKS!

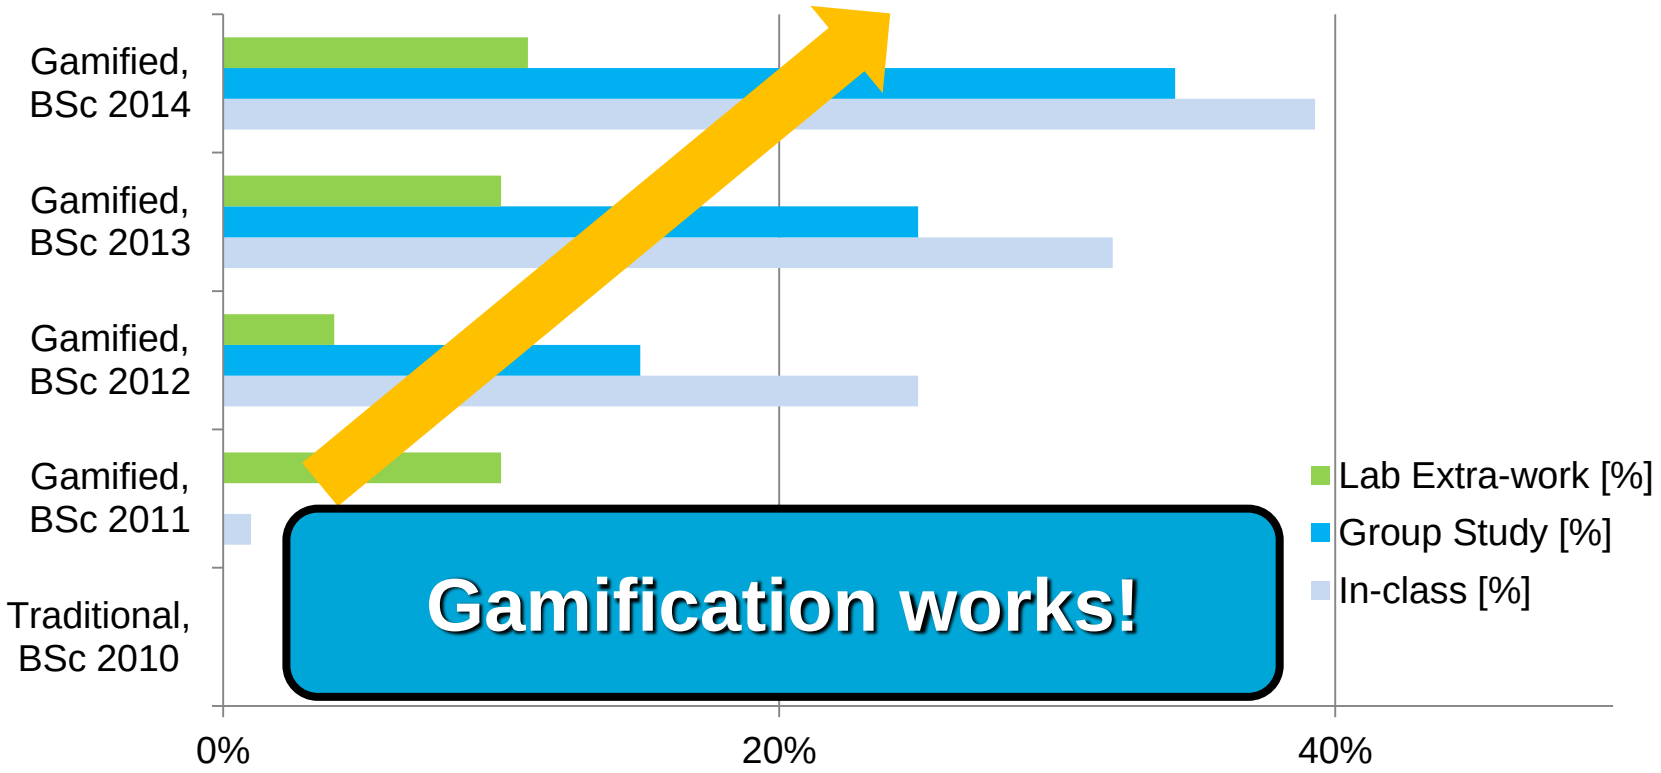

Extra work due to gamification, relative to traditional [% all students]

21ST CENTURY NEEDS FOR MASSIVIZING SYSTEMS

IS THIS SCIENCE TEACHABLE? DOING SCIENCE AS BSC + MSC STUDENT

IEEE Big Data 2013

IEEE/ACM SC15 Tutorial

PVLDB 2016

IEEE Scale Challenge Winners 2014

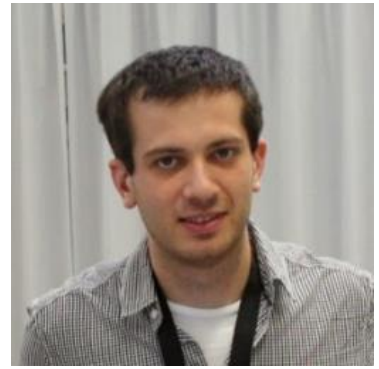

Bogdan Ghiț

VU

~40’

Massivizing Computer Systems
The Science of Distributed Ecosystems

~10’ — Vision, Current Golden Age of Ecosystems, and Crisis

~5’ — Intermezzo: What Is Modern Science?

~20’ — Massivizing Computer Systems

1. Stakeholders
2. Meaningful discovery: science + engineering + design
3. Experimental method for discovery + Reproducibility/Validation
4. Codified and teachable Science
5. What’s left?

~2’ — Take-home message

BUT WHAT IS "A SCIENCE"? WHAT IS A MODERN SCIENCE?

- Denning's criteria for a new field of science:
  - > Focus on pervasive phenomena
  - > By and for clear stakeholders
  - > Meaningful discovery
  - > Experimental methods of discovery
  - > Methods of validation
  - > Reproducibility of results
  - > Falsifiability of theoretical constructs
  - > Body of knowledge and skills that can be codified and taught

A SCIENCE OF ECOSYSTEMS

- My definition for a modern field of science:
  - > Fundamentals and Applications
  - > Methodology and Philosophy
  - > History and Envisioning
  - > Practice and Education
  - > Ethics

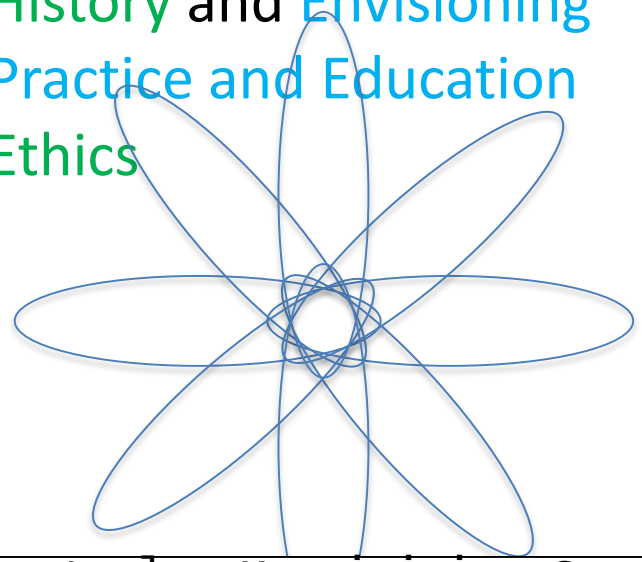

VU

Denning, The science in computer science, CACM 56, 2013.

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

MASSIVIZING COMPUTER SYSTEMS: WHAT'S LEFT?

A SCIENCE OF ECOSYSTEMS

THE VISION SEEMS GOOD, AND WE HAVE TAKEN PROMISING INITIAL STEPS

- Clients:
  - > Right to ICT
  - > Understanding

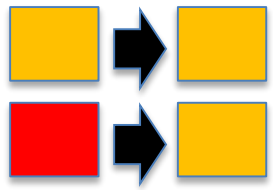

- Computer professionals:
  - > Create
  - > Operate
  - > Experiment
  - > Reproducible work
  - > Educate and train

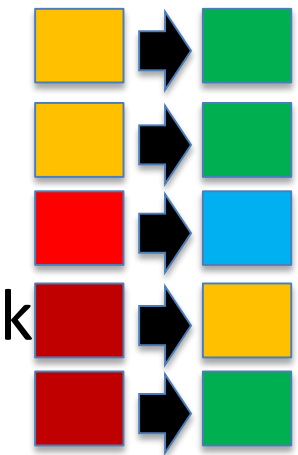

- ICT:
  - > Automated
  - > Efficient
  - > Controlled
  - > Ecosystems

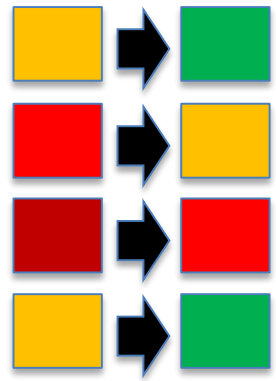

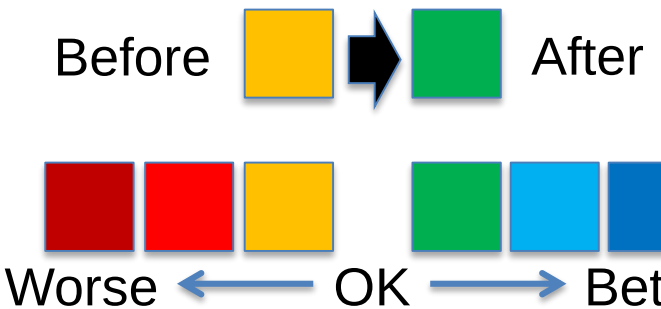

~40’

Massivizing Computer Systems
The Science of Distributed Ecosystems

~10’ — Vision, Current Golden Age of Ecosystems, and Crisis

~5’ — Intermezzo: What Is Modern Science?

~20’ — Massivizing Computer Systems

1. Stakeholders
2. Meaningful discovery: science + engineering + design
3. Experimental method for discovery + Reproducibility/Validation
4. Codified and teachable Science
5. What’s left?

~2’ — Take-home message

MASSIVIZING COMPUTER SYSTEMS

= THE SCIENCE OF MAKING COMPUTER SYSTEMS SCALABLE, RELIABLE, PERFORMANT, ETC., YET ABLE TO FORM EFFICIENT ECOSYSTEMS

- Golden Age of Computer Systems ... Yet crisis Is looming
- Massivizing Computer Systems
  - Think Ecosystems
  - Methods to address key challenges in science, engineering, and design
  - Teaching facilitated by award-winning method
- Much left to do, as we are merely beginning ...

@Large Research
Massivizing Computer Systems

http://atlarge.science

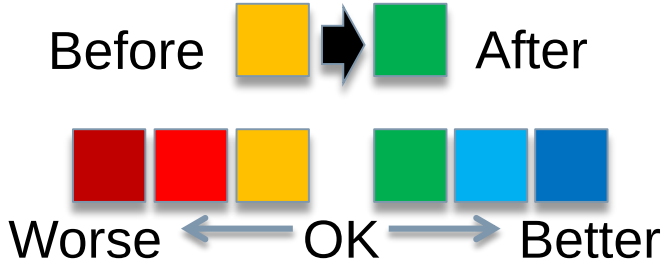

MASSIVIZING COMPUTER SYSTEMS

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

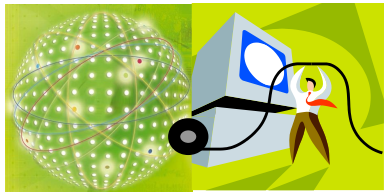

Collaboration or discussion about Massivizing Computer Systems:
Understanding, designing, deploying, tuning, analyzing, benchmarking distributed systems and ecosystems, including cloud computing and big data systems. Other topics in large-scale distributed systems and performance engineering are welcome.

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