

CLOUD COMPUTING (ECO)SYSTEMS

SCIENCE, DESIGN, AND ENGINEERING

Sponsored by:



<http://atlarge.science>

Many thanks to our collaborators.

Many thanks to our international working groups:



bit.ly/VUCloudEcosystems



@Alosup



Prof. dr. ir. Alexandru Iosup

VU AMSTERDAM < SCHIPHOL < THE NETHERLANDS < EUROPE



Amsterdam
founded 10th century
pop: 850,000









VU
founded 1880
pop: 23,500

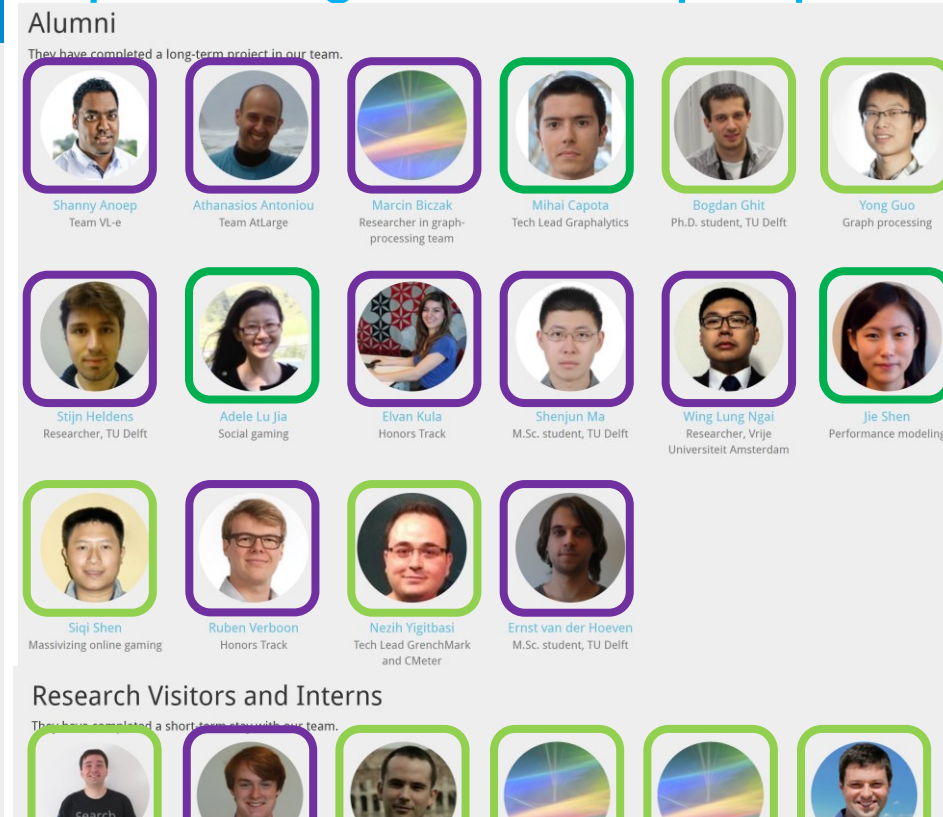


ATLARGE RESEARCH, OUR TEAM



<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



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.



THIS IS THE GOLDEN AGE OF CLOUD SYSTEMS AND ECOSYSTEMS




Education for
Everyone (Online)


Business
Services





Grid Computing



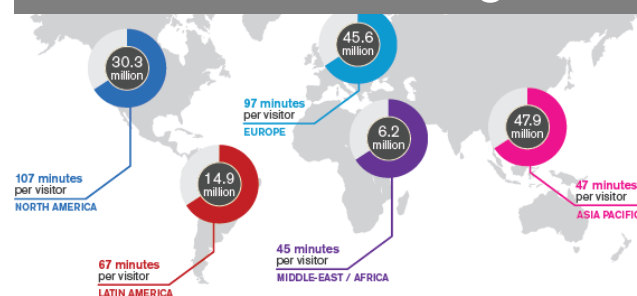
Grid
Computing



Big Science



Online Gaming



Datacenters



Daily Life

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


ONCE UPON A TIME ... THE DAWN OF THE CLOUD




Education for
Everyone (Online)



Business
Services



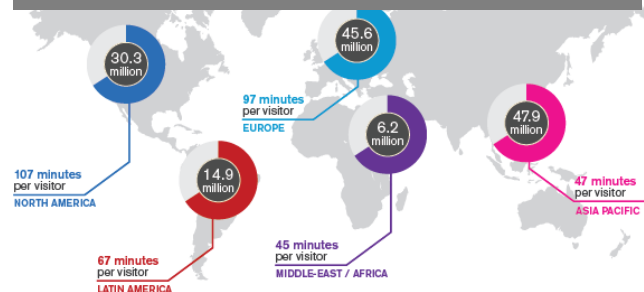
Big Science



CTAAGATGATCTTTAGTCCCGTTGAA
TCTTTAGTCCCGTTGATAACCAAC
GTAATACCAACCGGACTAAAGATCCG
GGACTAAAGTCCACCCCTATATATG

TTCAAAATTTCTCAAAAAGAGGGAG
GTGATTACATAAATCGGAGGTGCTA
TTGTCTACTACATTGACCTATGTTTT
GTAAGTTGATGAGAGAGAAATGTGT

Online Gaming



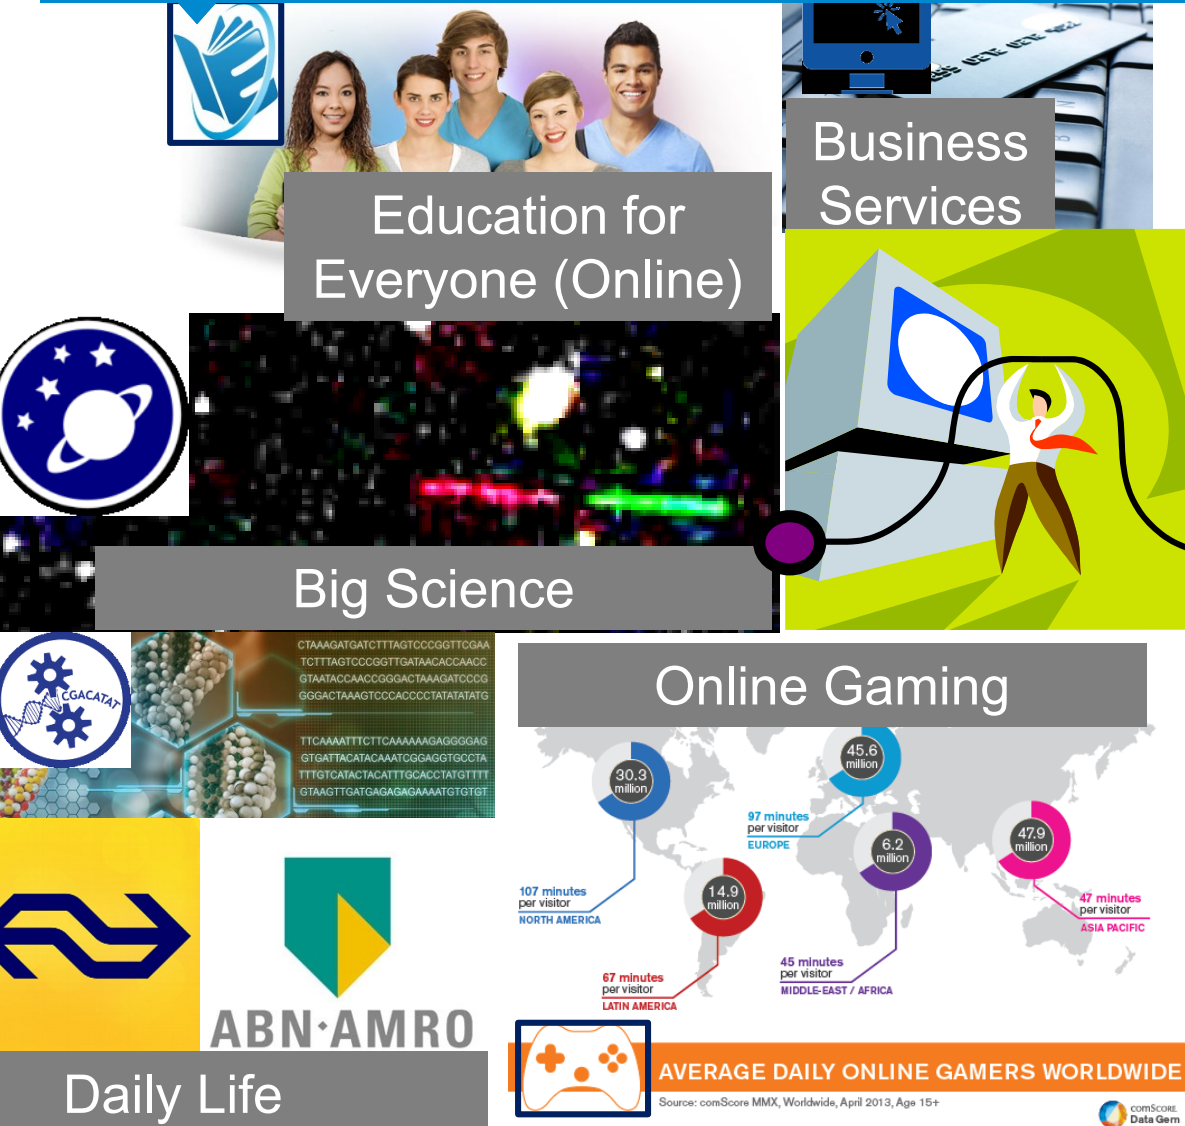
Daily Life

AVERAGE DAILY ONLINE GAMERS WORLDWIDE

Source: comScore MMX, Worldwide, April 2013, Age 15+




ONCE UPON A TIME ... THE DAWN OF THE CLOUD (1960s)



MIT Prof. Martin Greenberger:

“ Computing services and establishments will begin to spread throughout every life-sector [...] medical-information systems, [...] centralized traffic control, [...] catalogue shopping from [...] home, [...] integrated management-control systems for companies and factories [...] ”

M. Greenberger (1964) The Computers of Tomorrow
The Atlantic Monthly. Vol. 213(5), pp. 63-67, May.




ONCE UPON A TIME ... THE DAWN OF THE CLOUD (1960s)




Education for
Everyone (Online)



Business
Services



Big Science



Online Gaming




CTAAGATGATCTTTAGTCCCGTTTCGAA
TCTTTAGTCCCGTTGATAACCAACC
GTAATACCAACCGGACTAAAGATCCCG
GGGACTAAAGTCCCAACCCCTATATATG

TTCAAAATTTCTTCAAAAAGAGGGGAG
GTGATTACATACAAATCGGAGGTGCCTA
TTTGTCTACTACATTTGCACCTATGTTTT
GTAAGTTGATGAGAGAGAAATGTGTGT



Daily Life



Data Processing ~ SaaS
IBM-Service Bureau Corp.(SBC)

Software/System Dev. ~ PaaS
Computer Sciences Corp. (CSC)

Time Sharing ~ IaaS
IBM-SBC, Tymshare, GE Inf.Serv. (GEIS)

Facility management ~ IaaS
Electronic Data Systems (EDS)

Other Services
IBM

Source: J. R. Yost
(2017) Making IT Work.⁹

ONCE UPON A TIME ... THE DAWN OF THE CLOUD (1970s)



Education for
Everyone (Online)




Business
Services

Time Sharing ~ IaaS


GEIS invests in a large network

Tymshare invests in Tymnet

IBM invests in CALL 360



Big Science




Online Gaming

Technology
not ready

Emergence of
the PC

Slowdown:
1970 recession

DoJ anti-trust
lawsuit vs. IBM



AVERAGE DAILY ONLINE GAMERS WORLDWIDE

Source: comScore MMX, Worldwide, April 2013, Age 15+



Daily Life

Source: J. R. Yost
(2017) Making IT Work. ¹⁰

THIS IS THE GOLDEN AGE OF CLOUD COMPUTING (2010S)



Education for
Everyone (Online)




Business
Services



DIVERSE CLOUD SERVICES
FOR ALL

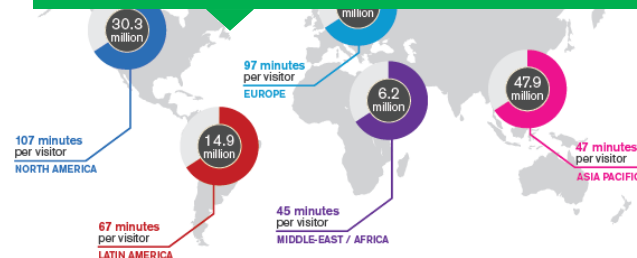
Grid Computing



Grid
Computing

Big Science

15-30% IT INDUSTRY



Datacenters



Daily Life

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

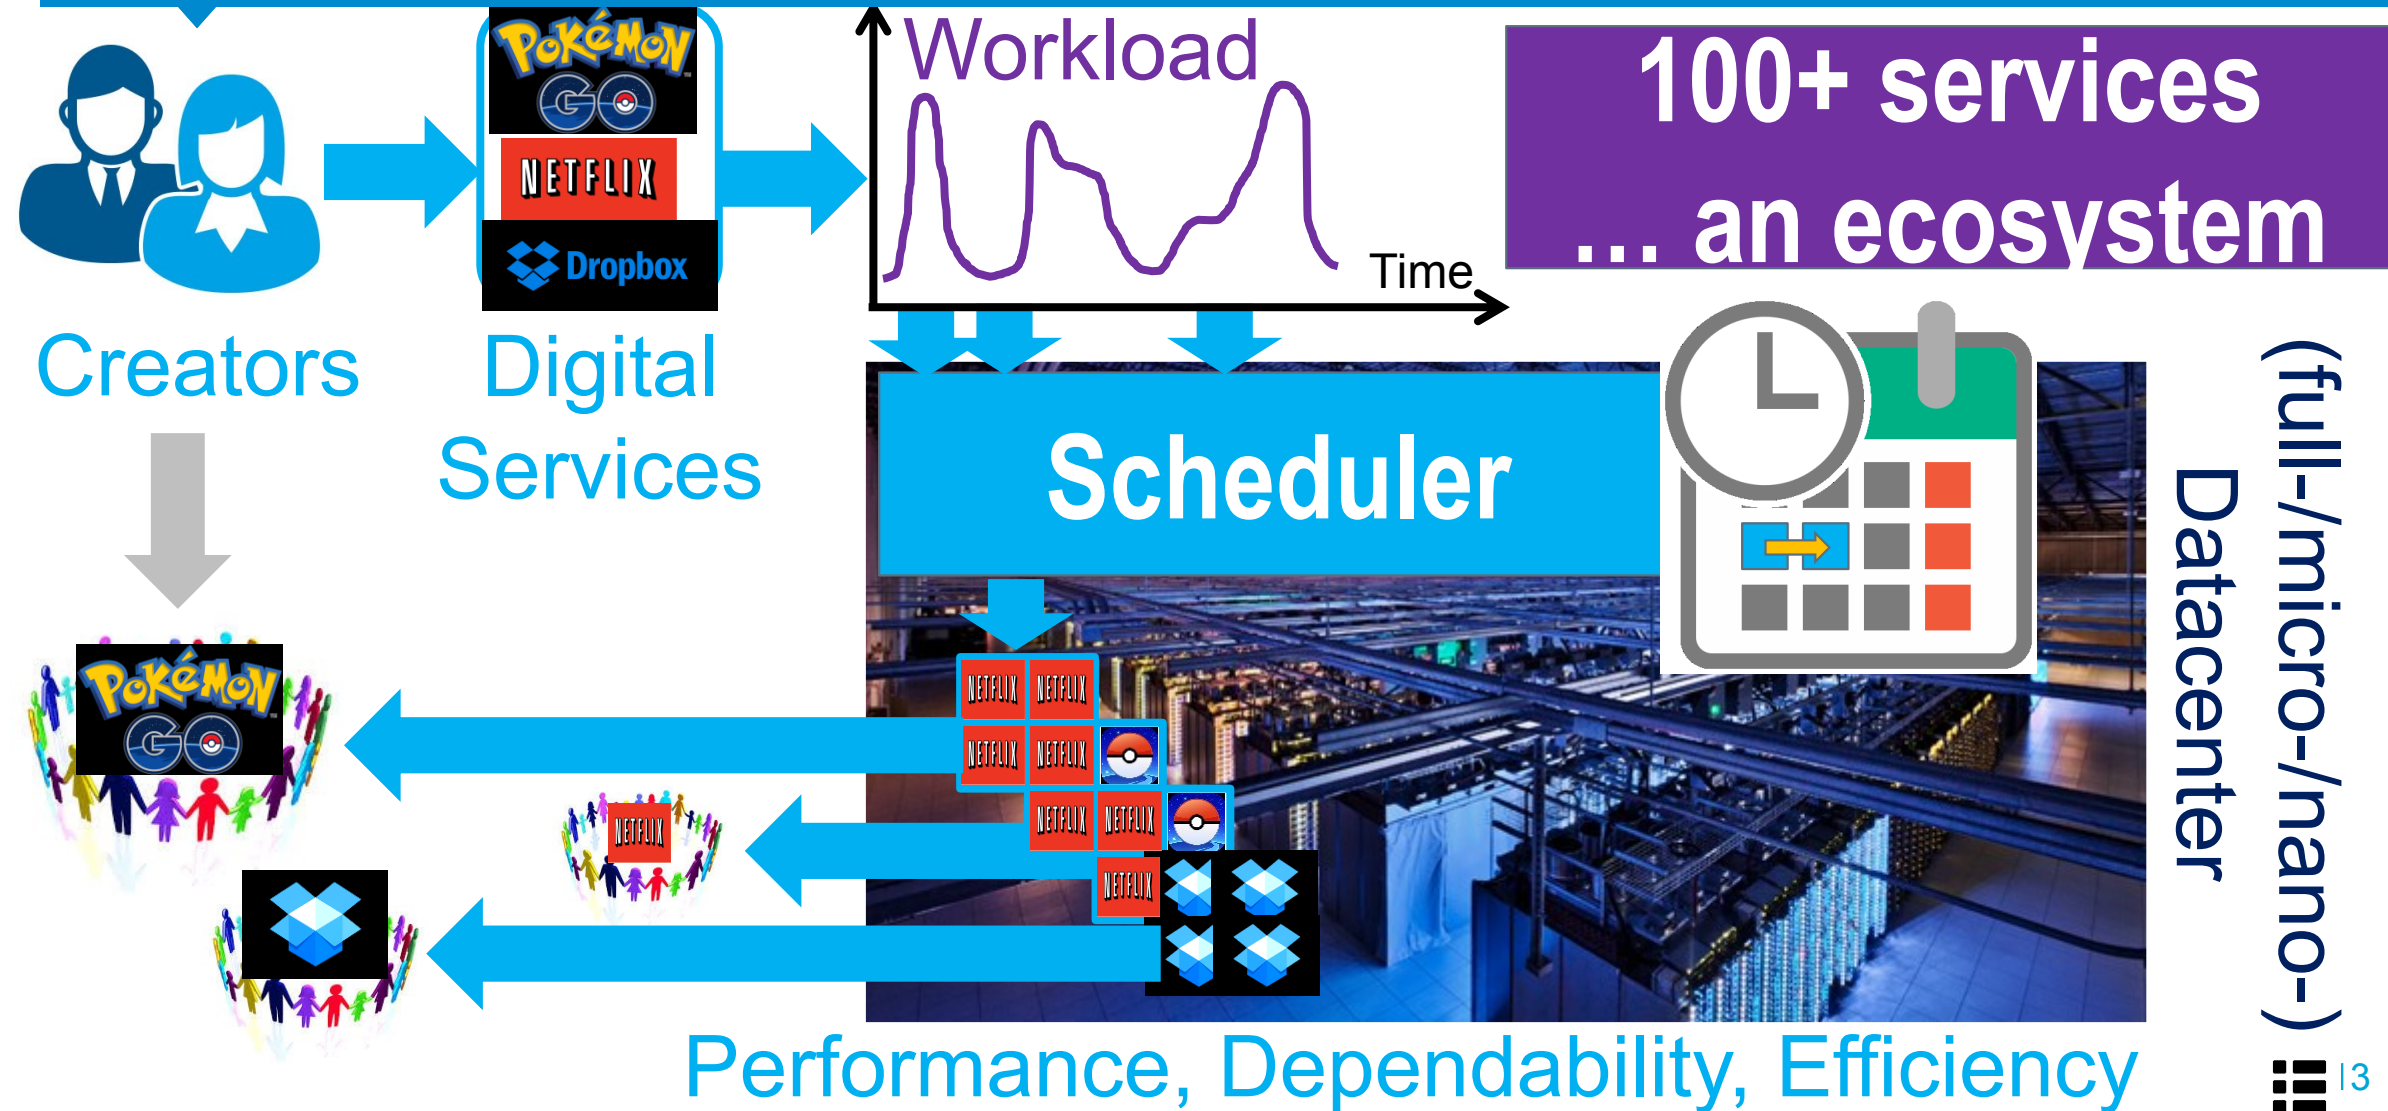
THIS IS THE GOLDEN AGE OF CLOUD COMPUTING (2010S)

Do you recognize this App?



Here is how it operates...

THE CLOUD ECOSYSTEM: SERVICE, DATACENTER, SCHEDULER



DIVERSE CLOUD SERVICES FOR ALL... ARE WE THERE YET?



Technology
not ready


Emergence of
the mobile... 5G

Anti-trust
lawsuits vs. GAFAM

Not covered in this talk



Technology
not ready



Pokémon GO Server Status

REFRESH

Pokémon GO

OFFLINE
for 15 minutes

Pokémon Trainer Club

UNSTABLE
for 2 minutes

Pokémon GO Uptime

55.56%
over the past hour


96.29%
over the past day

Pokémon Trainer Club Uptime

66.67%
over the past hour

96.66%
over the past day

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



Pokémon GO Server Status

REFRESH

Pokémon GO

OFFLINE
for 15 minutes

Pokémon Trainer Club

UNSTABLE
for 2 minutes

Pokémon GO Uptime

55.56%
over the past hour

96.29%
over the past day

Pokémon Trainer Club Uptime

66.67%
over the past hour

96.66%
over the past day

My Research: Massivizing Computer Systems

Technology not ready

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.

science + engineering + design

REPRODUCIBILITY AND VALIDATION OF DISCOVERY

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

METHODOLOGY

OPEN SCIENCE

REPORTING &
DISSEMINATION

REPRODUCIBILITY

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

Munafò et al., A manifesto for reproducible science, Nature Human Behaviour, Jan 2017. [[online](#)]

REPRODUCIBILITY AND VALIDATION OF DISCOVERY

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

METHODOLOGY

OPEN SCIENCE

REPORTING &
DISSEMINATION

REPRODUCIBILITY



A. V. PAPADOPOULOS, L. VERSLUIS, A. BAUER, N. HERBST, J. VON KISTOWSKI, A. ALI-ELDIN, C. ABAD, J. N. AMARAL, P. TUMA, AND A. IOSUP (2019) METHODOLOGICAL PRINCIPLES FOR REPRODUCIBLE PERFORMANCE EVALUATION IN CLOUD COMPUTING.

IEEE TRANSACTIONS ON SOFTWARE ENGINEERING 2019 (IN PRINT)

SPEC RG CLOUD GROUP. TECHNICAL REPORT. APR 2019 [[ONLINE](#)]

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

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

THE DESIGN OF DISTRIBUTED SYSTEMS AND ECOSYSTEMS

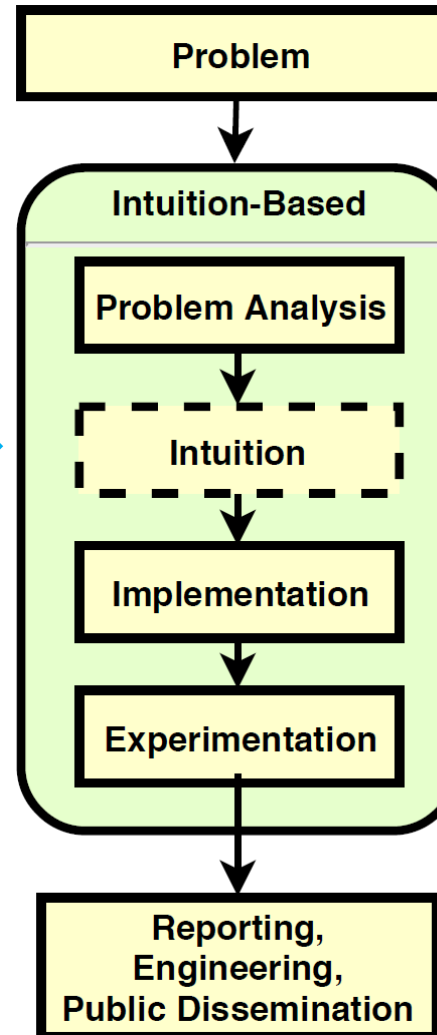

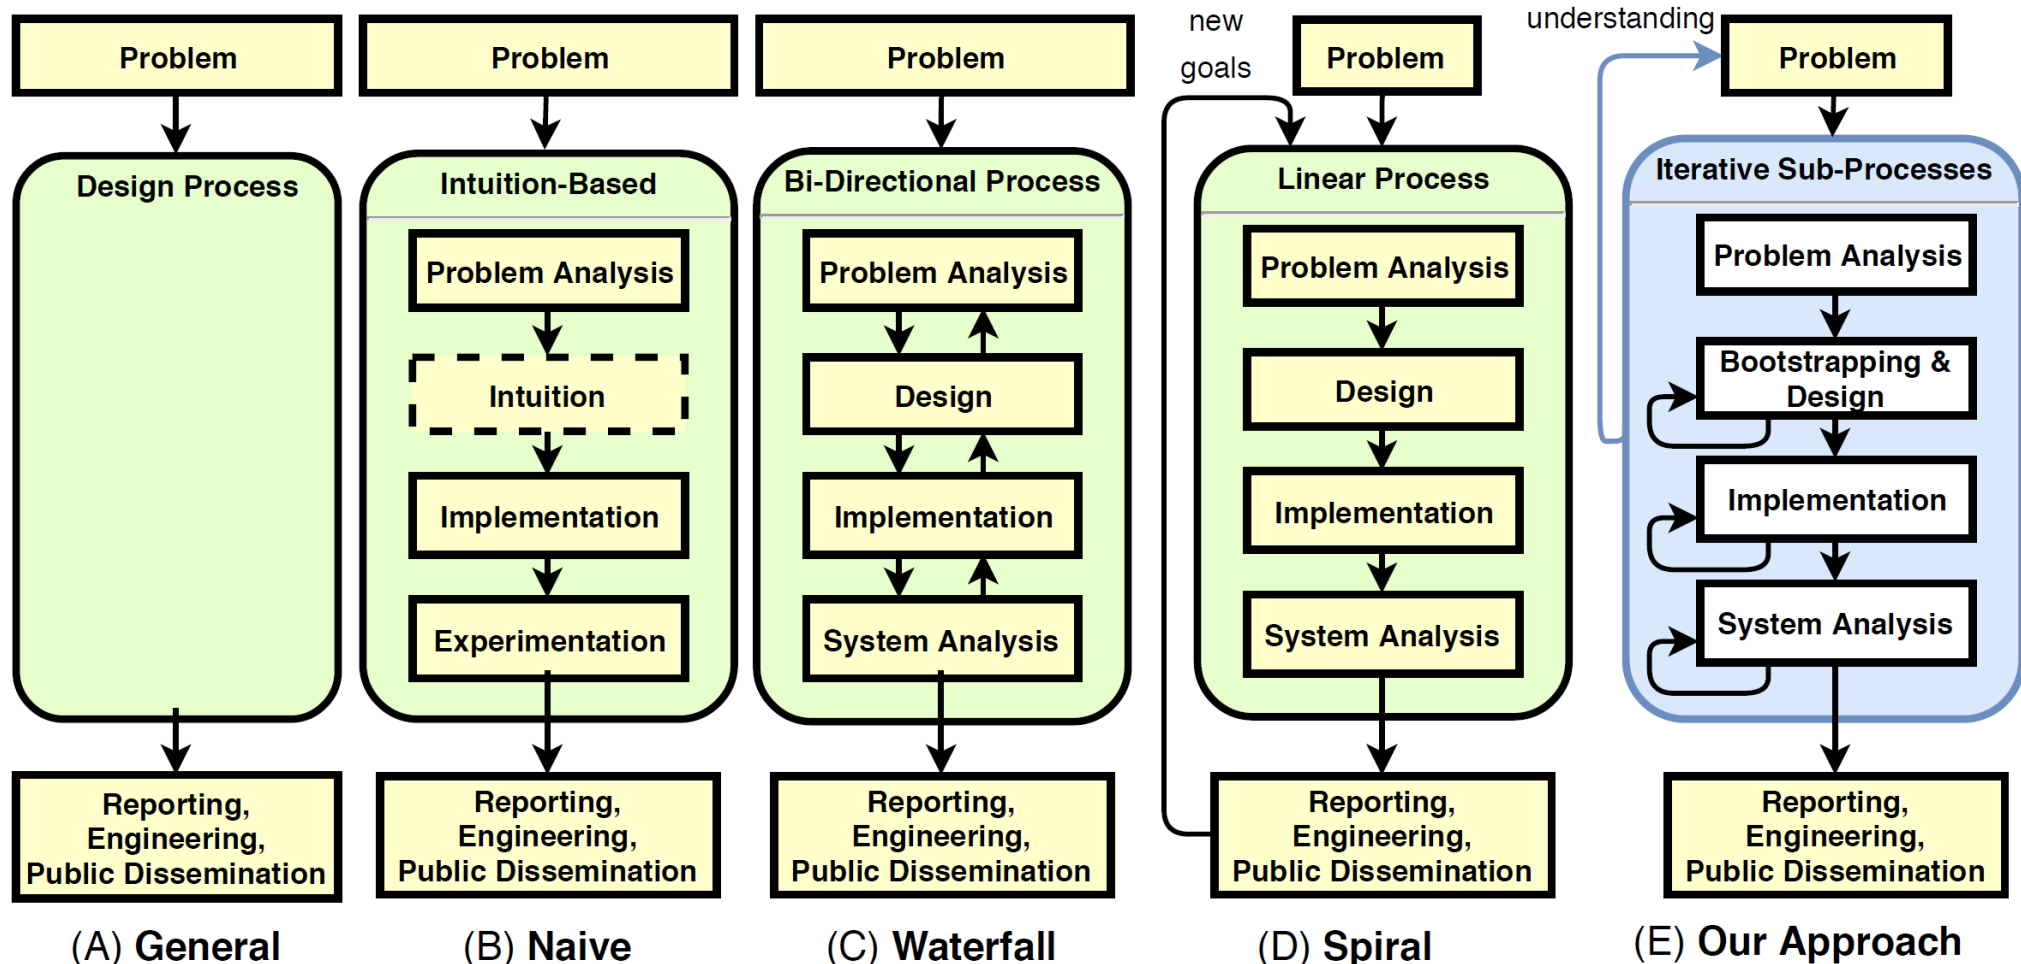


Photo by Matthew Yohe, 2008.

CC 3.0 Some rights reserved.

THE ATLARGE DESIGN PROCESS FOR DISTRIBUTED SYSTEMS AND ECOSYSTEMS

bit.ly/AtLargeDesign1Talk




[Iosup et al. The AtLarge Vision on the Design of Distributed Systems and Ecosystems. ICDCS'19] online version: <http://arxiv.org/pdf/1902.05416>

CHALLENGE: MEET SERVICE LEVEL AGREEMENTS

PERFORMANCE, DEPENDABILITY, AND OTHER NON-FUNCTIONAL CHALLENGES

We Cannot Maintain the Ecosystems we Have Built (and Thought We've Tested, and Validated)



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

Goog
world
System

Cloudflare and Google dealt with issues that affected countless sites and users on Tuesday.

By David Yaff

July 2, 2019



When a website won't load, many internet users turn to DownDetector, a site that keeps track of online disruptions, providing frequent updates infrastructure.

Sources: <https://www.nytimes.com/2019/07/02/business/cloudflare-google-internet-problems.html>

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](#)]

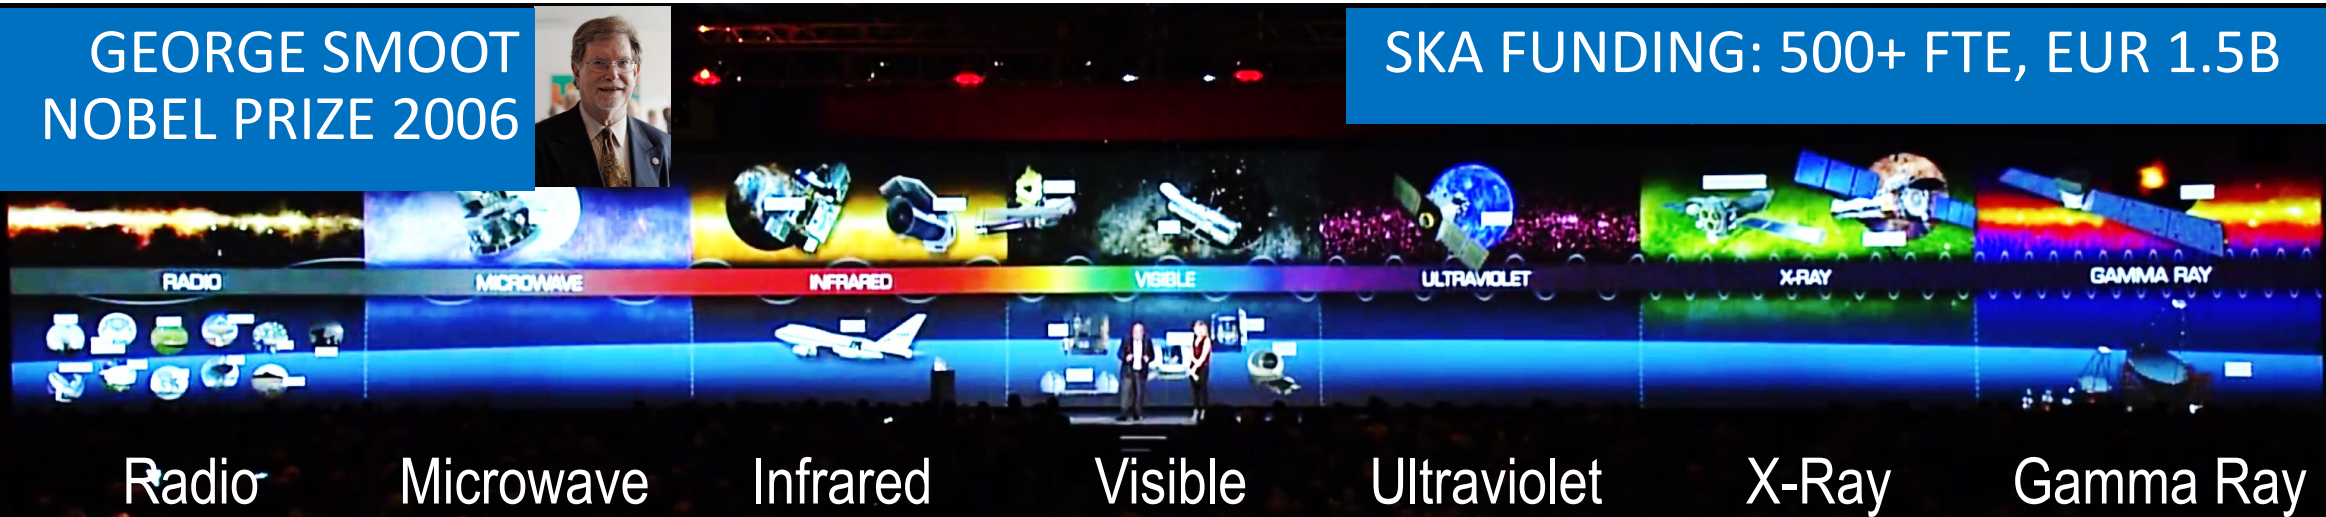
DISCOVERY = LARGE-SCALE, LONG-TERM STUDY

UNCOVERING THE MYSTERIES OF OUR PHYSICAL UNIVERSE

GEORGE SMOOT
NOBEL PRIZE 2006



SKA FUNDING: 500+ FTE, EUR 1.5B



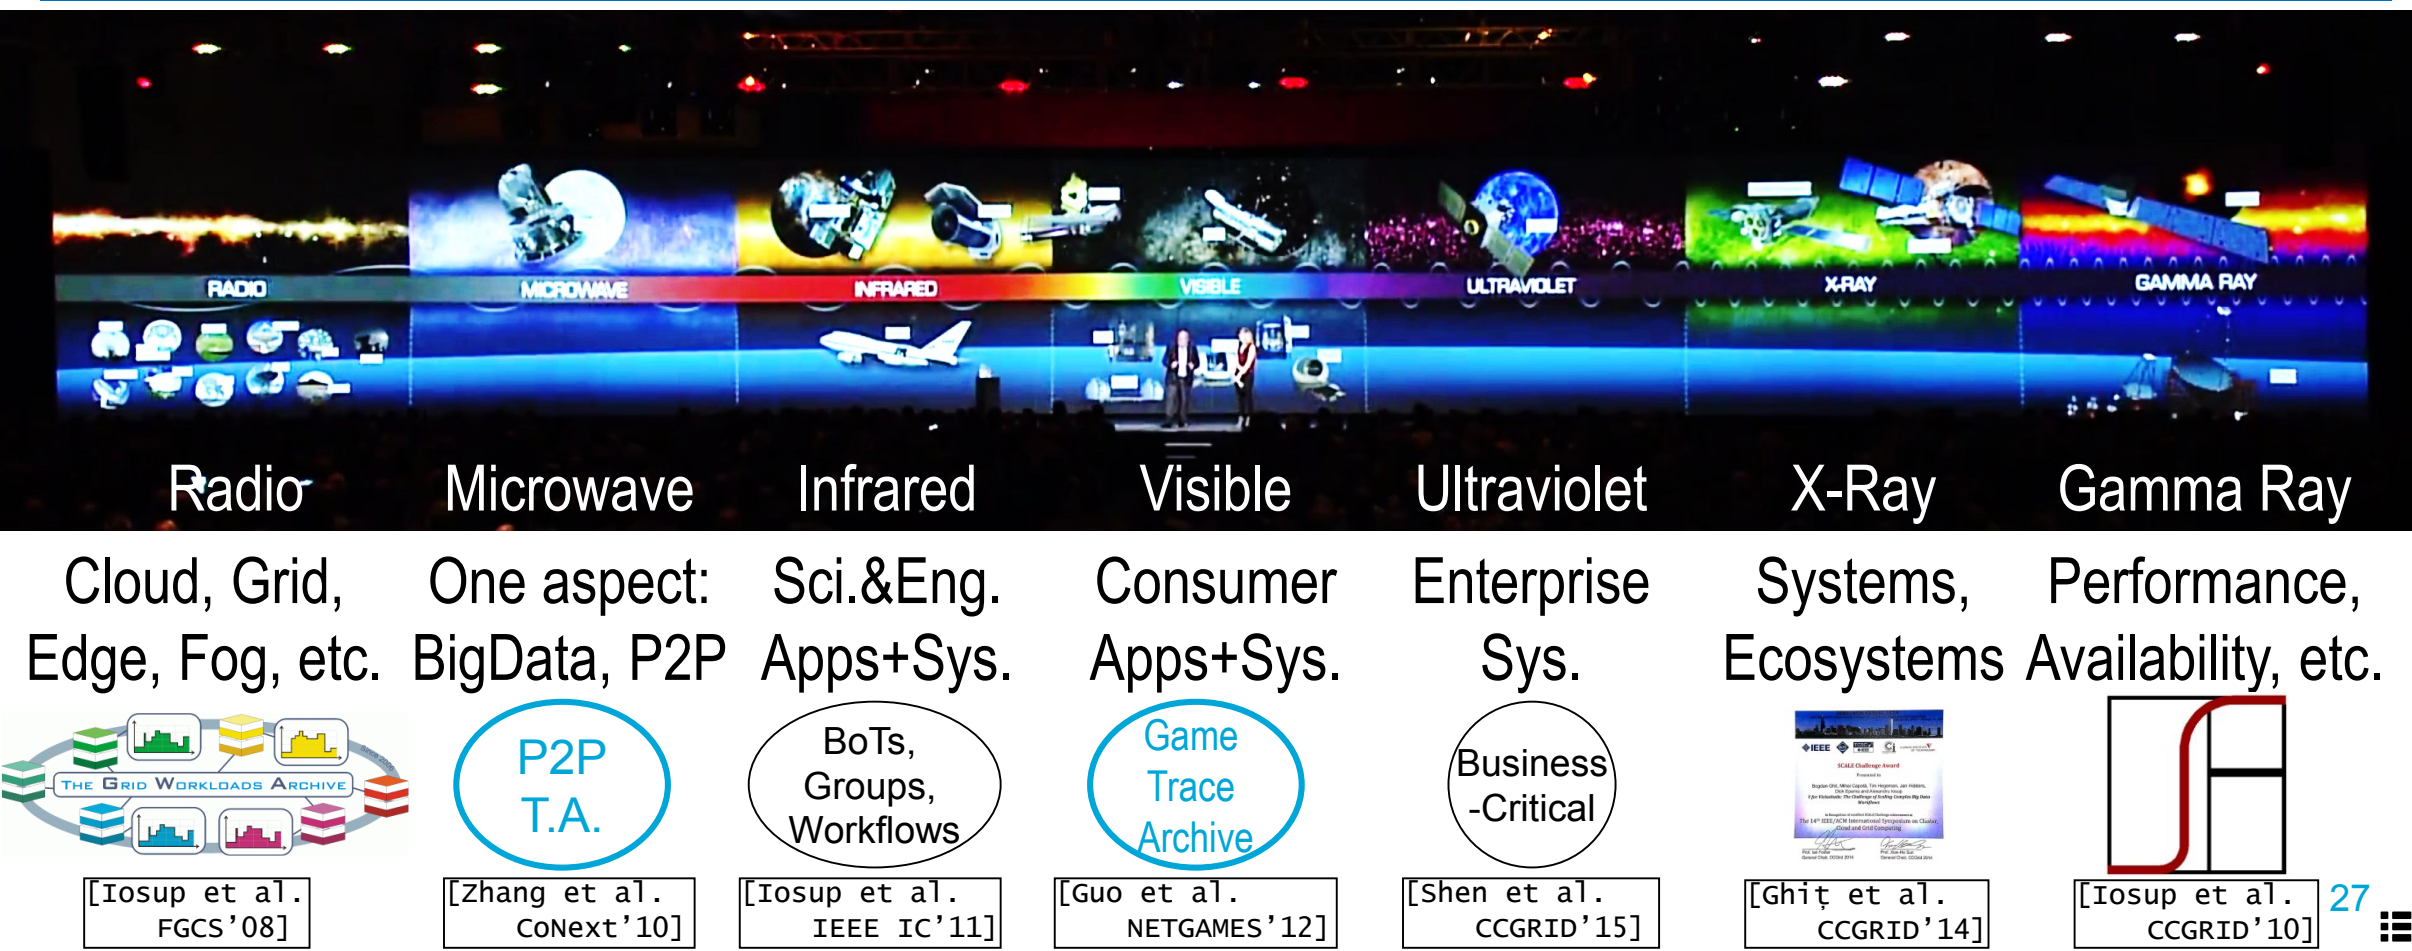
James Cordes, The Square Kilometer Array, Project Description, 2009 [[online](#)]

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



THE WORKFLOW TRACE ARCHIVE

96 traces






Laurens
Versluis

METADATA AND TRACES FOR YOUR WORKFLOW SYSTEMS

WORKFLOWS ARE COMMON
IN MANY DOMAINS

EXCEPT IN SCI.,
DESIGN, & ENG.

THE WORKFLOW TRACE ARCHIVE
CORRECTS THIS



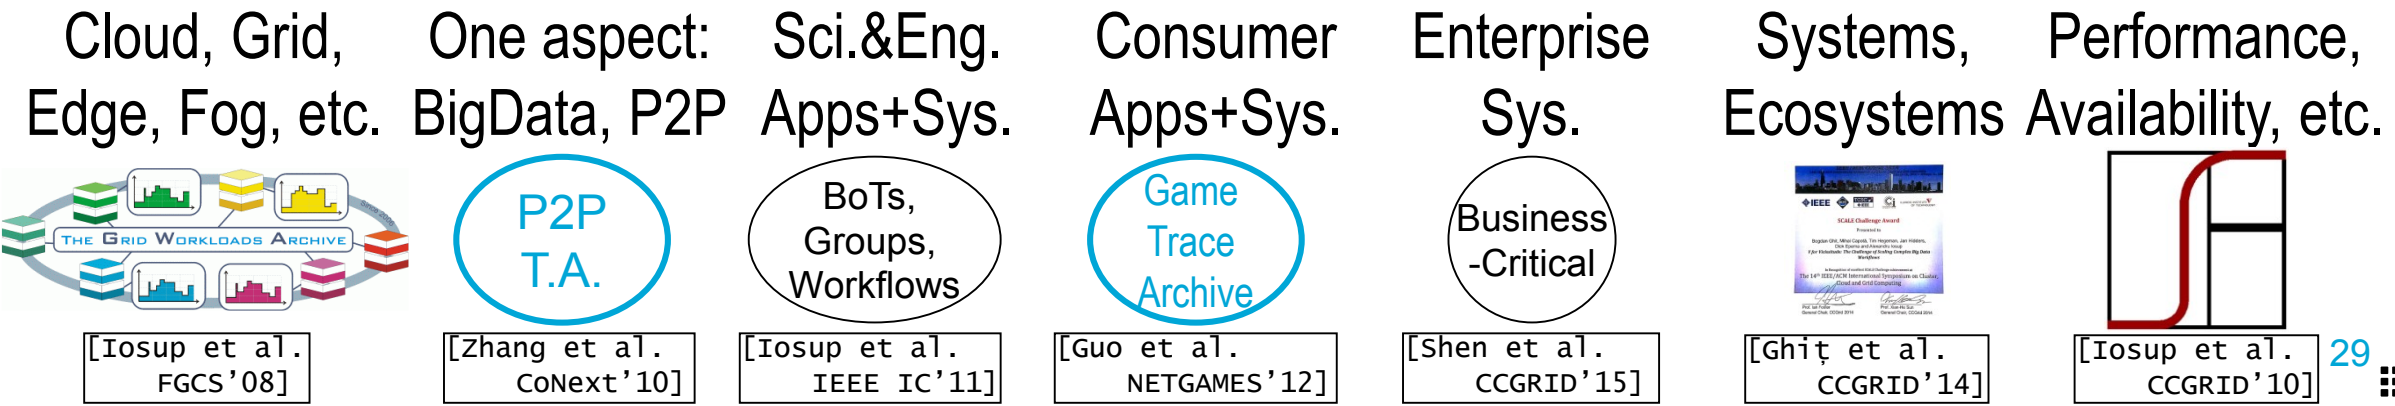
[Versluis et al. The workflow Trace Archive]
Tech.rep.: <http://arxiv.org/pdf/1906.07471>

[http:// wta.atlarge.science](http://wta.atlarge.science)

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!



UNKNOWN PHENOMENA: INTER-, ADAPT-, EXAPTATION

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

Toward a Distributed Systems Memex



...

...

...

...

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, Performance,
Ecosystems 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]



[Ghiț et al.
CCGRID'14]



[Iosup et al.
CCGRID'10]




EXPERIMENTAL METHODS OF DISCOVERY

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


Our Prototypes (*in physico/in vitro*)



SURFnet6



Alex Uta



Georgios
Andreadis




Fab
Mastel



Maria Voinea



Laurens
Versluis



Alexey Ilyushkin



We also use clouds



OpenDC

And simulators (*in silico*)




EXPERIMENTAL METHODS OF DISCOVERY

UNDERSTANDING THE PHENOMENON OF PERFORMANCE VARIABILITY



Our Prototypes (*in physico/in vitro*)



SURFnet6



Alex Uta



EXPERIMENTAL METHODS OF DISCOVERY

NETWORK VARIABILITY AFFECTS BIG DATA PERFORMANCE


→ NEW DESIGNS NEEDED

Our Prototypes (*in physico/in vitro*)



Alex Uta

Application	Maximum Slowdown	Bandwidth Distribution
Wordcount	1.61	A
Sort	1.51	D
Terasort	1.79	A
K-Means	1.48	D
Bayes	1.14	A
Pagerank	1.07	A




CHALLENGE: SYSTEMATIC DESIGN & DESIGN-SPACE EXPLORATION

THE COMPLEXITY CHALLENGE



We Build and Test
Isolated Computer Systems, Yet
Everything Works in Stacks and Ecosystems
... Need to Reason About Them!

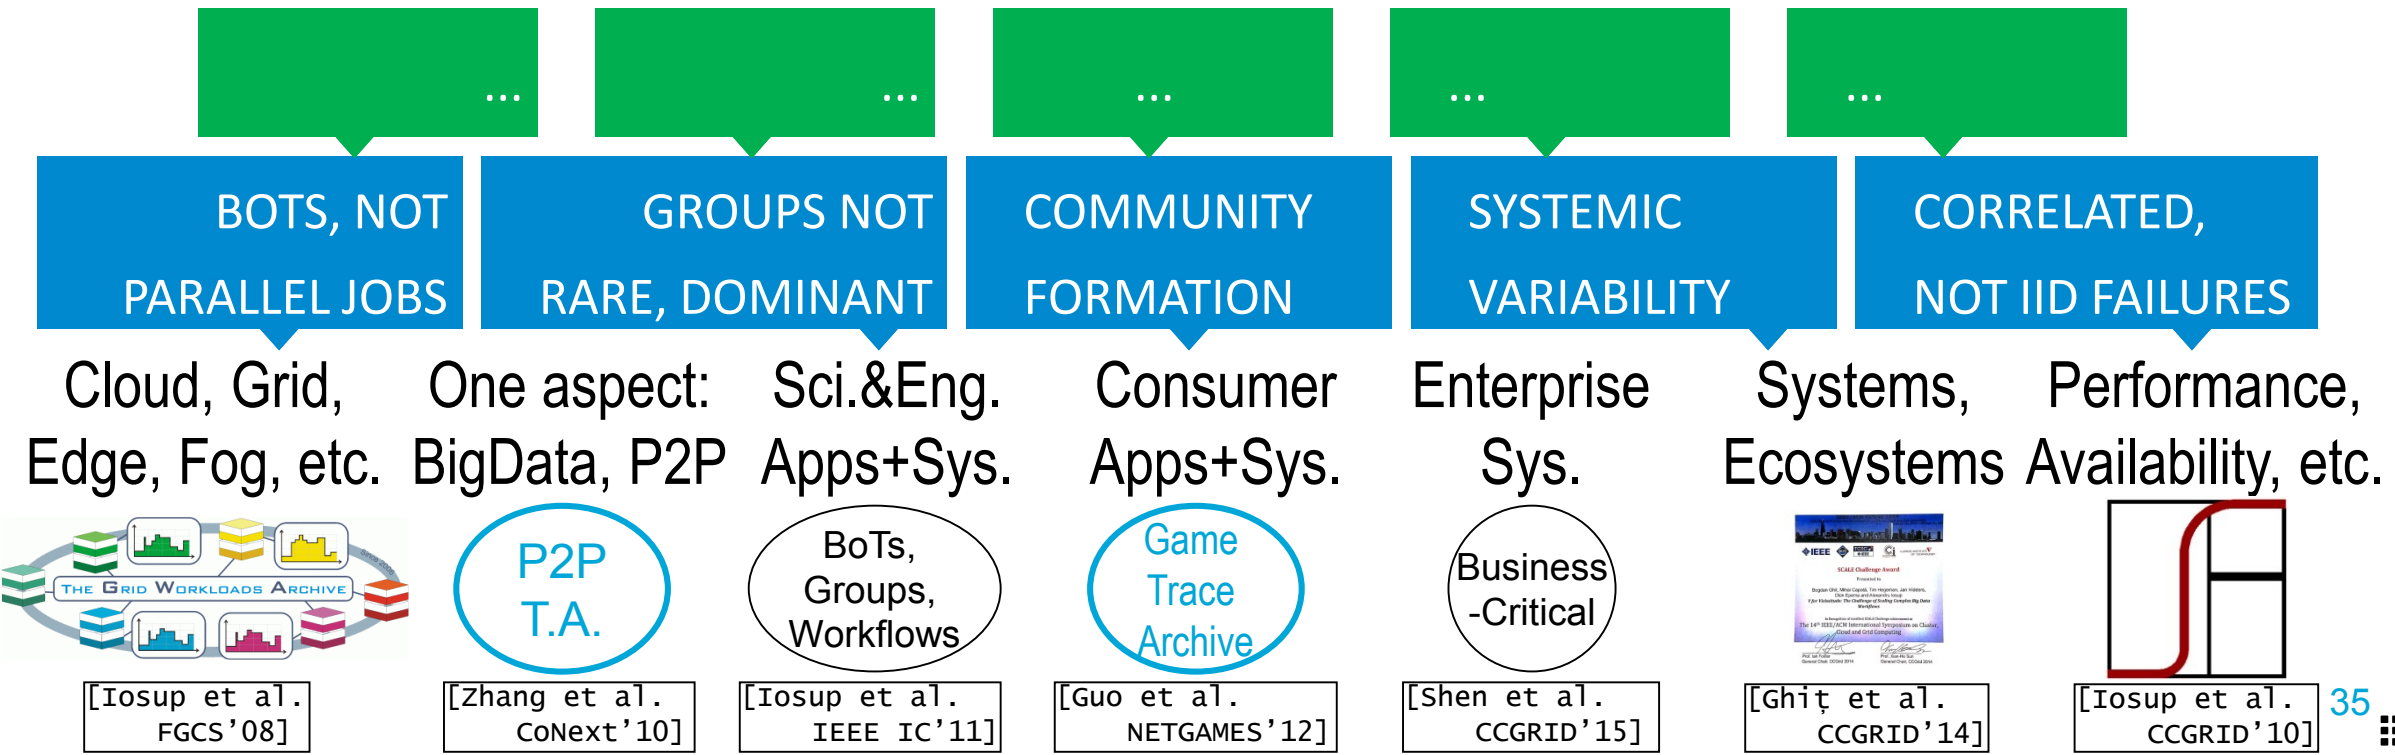


<<1% OF BIG DATA BY MATT TURK (2017)
“SW. IS EATING THE WORLD”

UNKNOWN PHENOMENA: INTER-, ADAPT-, EXAPTATION

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

BUT ... HOW CAN DESIGNERS MANAGE COMPLEXITY, WITH SO MANY PHENOMENA?



MEANINGFUL DISCOVERY

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



- 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

FUNDING: > 3B USD

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. [\[Online\]](#)

HOW TO MANAGE SYSTEM COMPLEXITY?

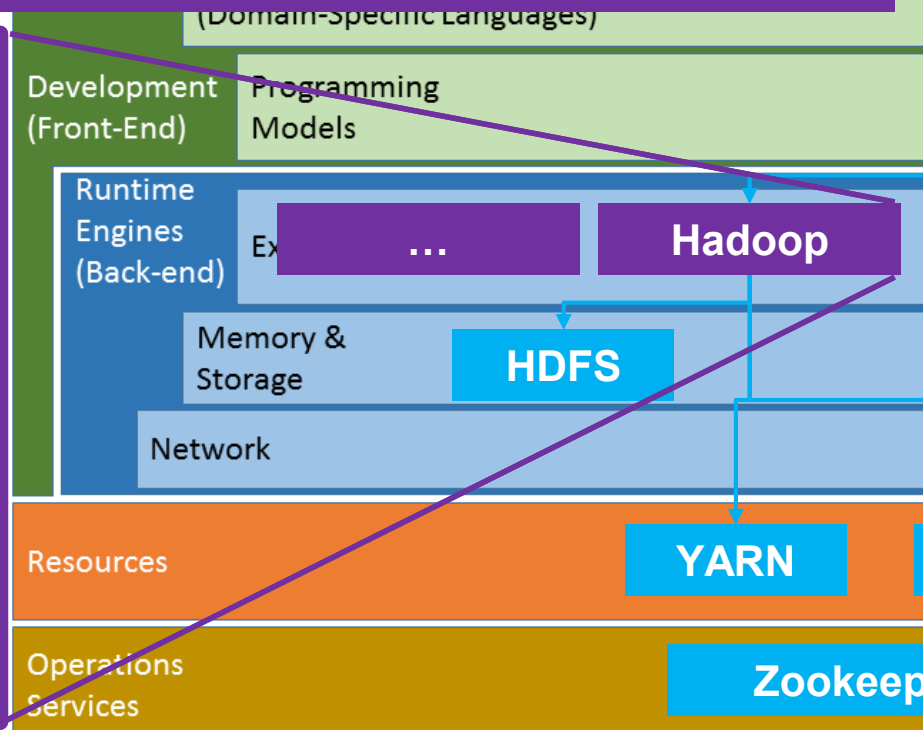
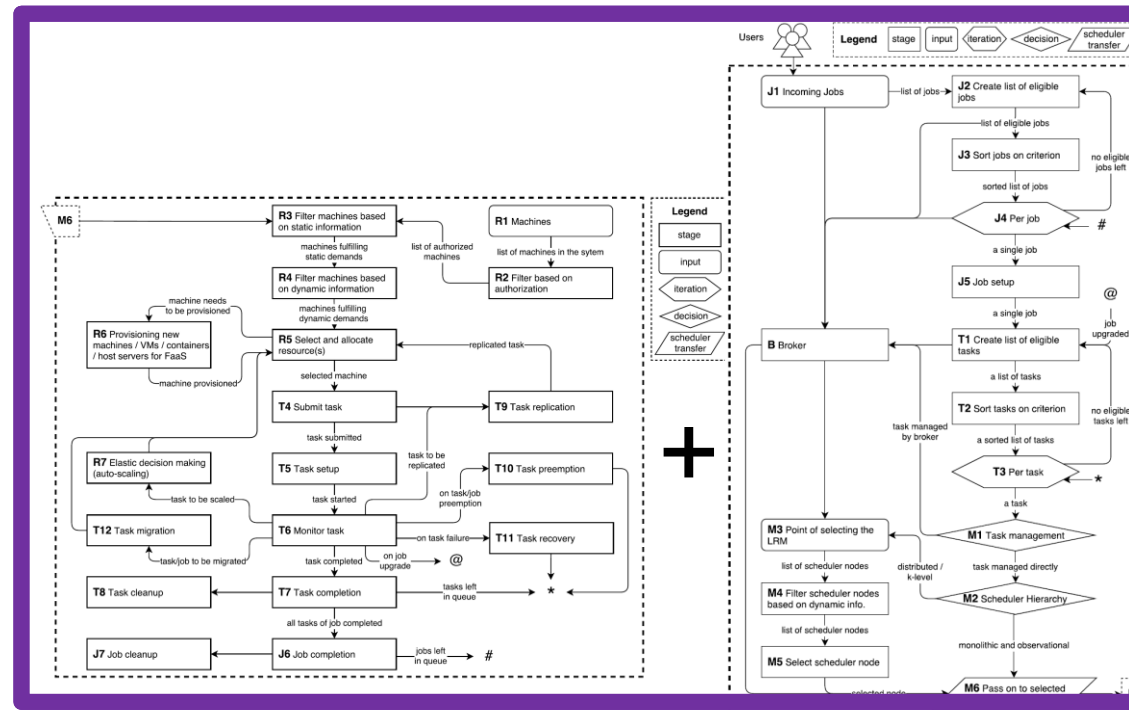
THE COMPLEXITY CHALLENGE

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS



Georgios Andreadis

ANDREADIS ET AL. REFERENCE ARCHITECTURE FOR SCHEDULERS IN DCS



[Andreadis et al. SC'18]

1

Infrastructure Physical

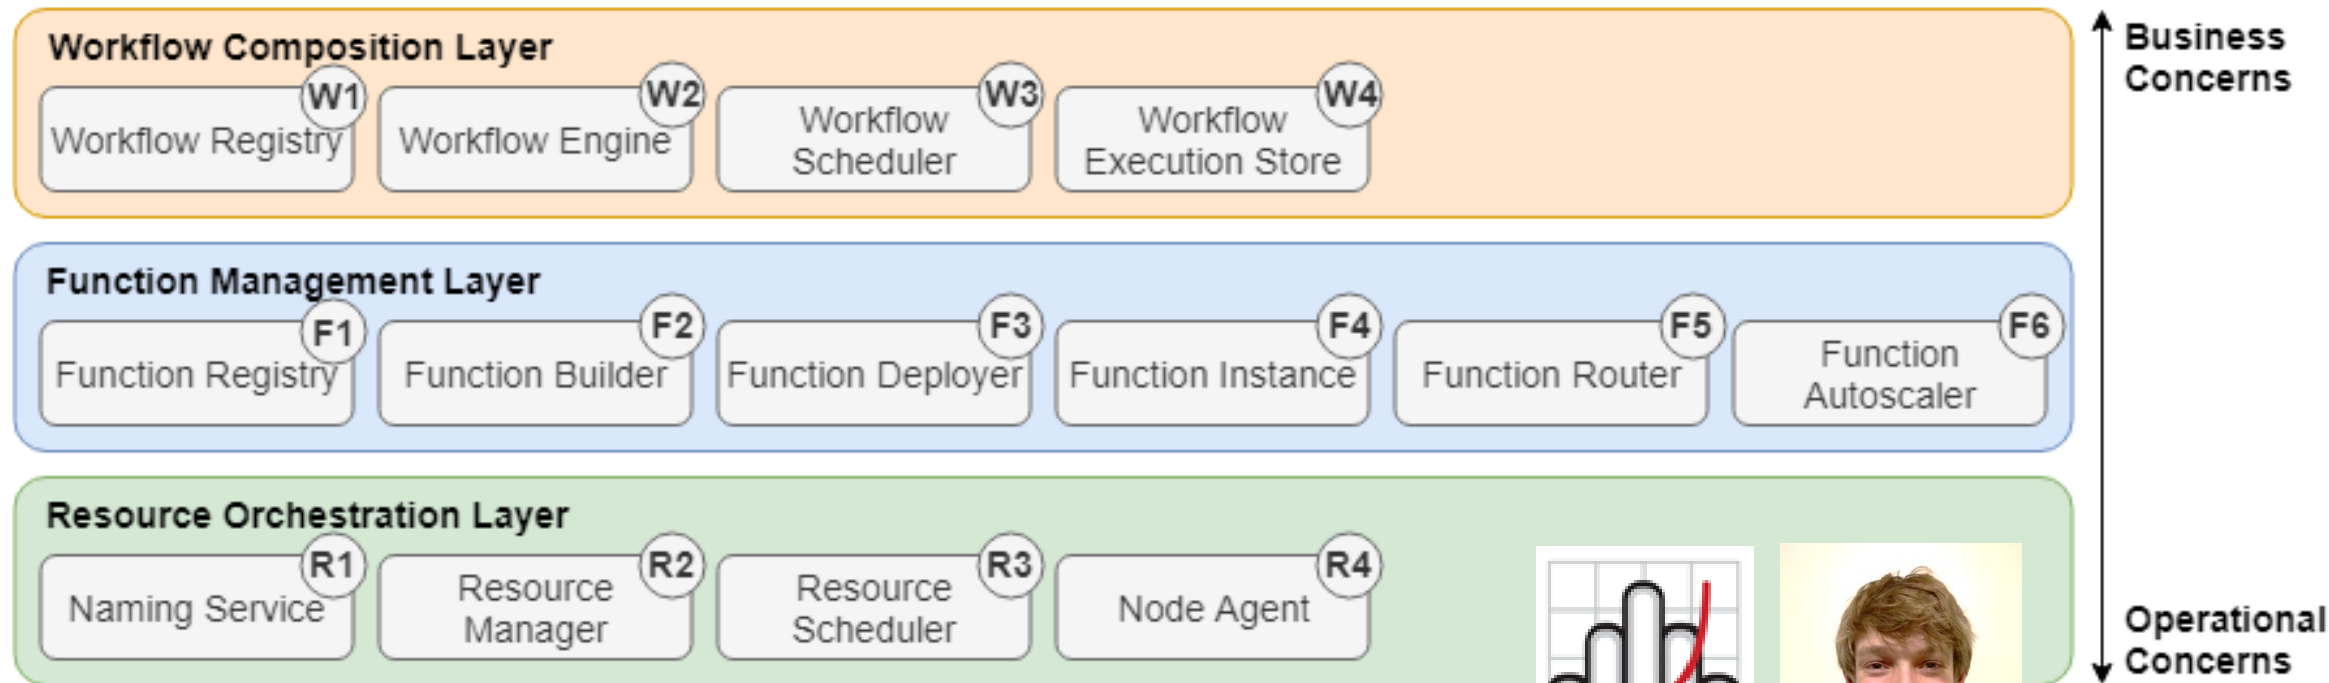
Physical Architecture/Hierarchy: DC, Room/Container, Pod/Partition, Cluster, Rack

Node


HOW TO MANAGE SYSTEM COMPLEXITY?

THE COMPLEXITY CHALLENGE

WIP REFERENCE ARCHITECTURE OF FAAS PLATFORMS



[van Eyk et al. (2018) Serverless is More: From PaaS to Present Cloud Computing, IEEE Internet Computing] [[online](#)]



Erwin
van Eyk

CHALLENGE: EFFICIENCY, SUSTAINABILITY, RESPONSIBILITY!

THE RESOURCE MANAGEMENT CHALLENGE

Based on Jav Walker's recent TED talk.

Need To Be Much More Efficient,

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).



Erwin
van Eyk



Alexandru
Iosup



Serverless / FaaS Execution

Vision and Architecture for Serverless Execution in Cloud Environments



van Eyk, Toader, Talluri, Versluis, Uta, Iosup: Serverless is More: From PaaS to Present Cloud Computing. IEEE Internet Computing Sep/Oct 2018. [[Online](#)]

Erwin Van Eyk, Alexandru Iosup, Cristina L. Abad, Johannes Grohmann, Simon Eismann: A SPEC RG Cloud Group's vision on the Performance Challenges of FaaS Cloud Architectures. ICPE 2018. [[Online](#)]


Erwin van Eyk, Simon Seif (SAP), Markus Thoemmes (IBM Germany), Alexandru Iosup. The SPEC Cloud Group's Research Vision on FaaS and Serverless Architectures. workshop on Serverless Computing (woSC'17), held in conjunction with Middleware'17. [[Online](#)]

SERVERLESS STREAMING WORKFLOWS

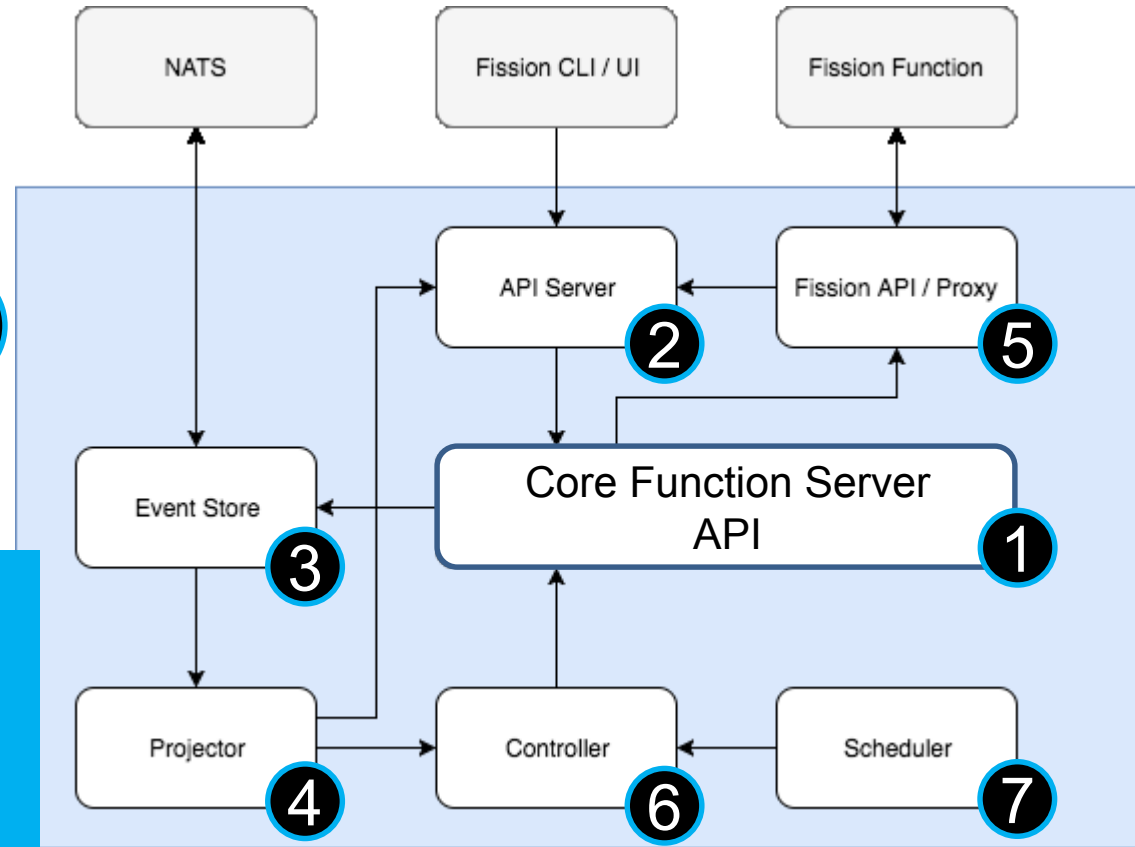
DESIGNING SERVERLESS ARCHITECTURES, APIS, AND SCHEDULERS




Erwin
van Eyk



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



MASSIVIZING COMPUTER SYSTEMS



FURTHER READING

<https://atlarge-research.com/publications.html>

1. Iosup et al. Massivizing Computer Systems. ICDCS 2018 ← start here
2. Iosup et al. The AtLarge Vision on the Design of Distributed Systems and Ecosystems. ICDCS 2019
3. Papadopoulos et al. Methodological Principles for Reproducible Performance Evaluation in Cloud Computing. IEEE Trans. on Sw. Eng. 2019
4. Van Eyk et al. Serverless is More: From PaaS to Present Cloud Computing, IEEE IC Sep/Oct 2018
5. Andreadis et al. A Reference Architecture for Datacenter Scheduling, SC18
6. Talluri et al. Characterization of a Big Data Storage Workload in the Cloud. ACM/SPEC ICPE 2019.
7. Ilyushkin et al. Autoscalers. TOMPECS 2018.
8. Uta et al. Exploring HPC and Big Data Convergence.
9. Uta et al. Elasticity in Graph Analytics? IEEE Cluster 2018.
10. Toader et al. Graphless. IEEE ISPDC'19.
11. Herbst et al. Ready for rain? TOMPECS 2018.
12. Guo et al. Streaming Graph-partitioning. JPDC'18.
13. Iosup et al. The OpenDC Vision. ISPDC'17.
14. van Beek et al. Managing Operational and Disaster-Recovery Risks in Virtualized DCs. ISPDC'19.
15. Iosup et al. LDBC Graphalytics. PVLDB 2016.

CLOUD SYSTEMS AND ECOSYSTEMS

PART OF THE LARGER VISION OF MASSIVIZING COMPUTER SYSTEMS



Many thanks to
200+
collaborators

- Golden Age of Cloud Ecosystems ... Yet many challenges
 1. Reproducibility
 2. Design Process
 3. Trace Archives
 4. Phenomena:
performance variability, etc.
 5. Reference Architecture
 6. Serverless engine
 7. Benchmarking (*)
 8. Simulation for
(datacenter) operation (*)
(*) extra slides
- Much left to do ... Time to engage with the SPEC RG Cloud Group!

