

MASSIVIZING COMPUTER SYSTEMS

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

@Large Research
Massivizing Computer Systems



<http://atlarge.science>


Co-sponsored by:



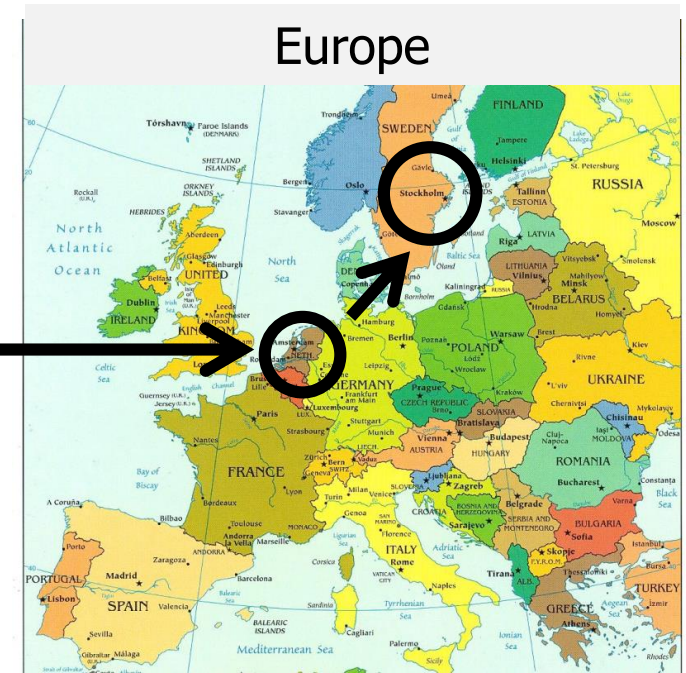
VU AMSTERDAM < SCHIPHOL < THE NETHERLANDS < EUROPE



Amsterdam
founded 10th century
pop: 850,000



VU
founded 1880
pop: 23,500



WHO AM I?

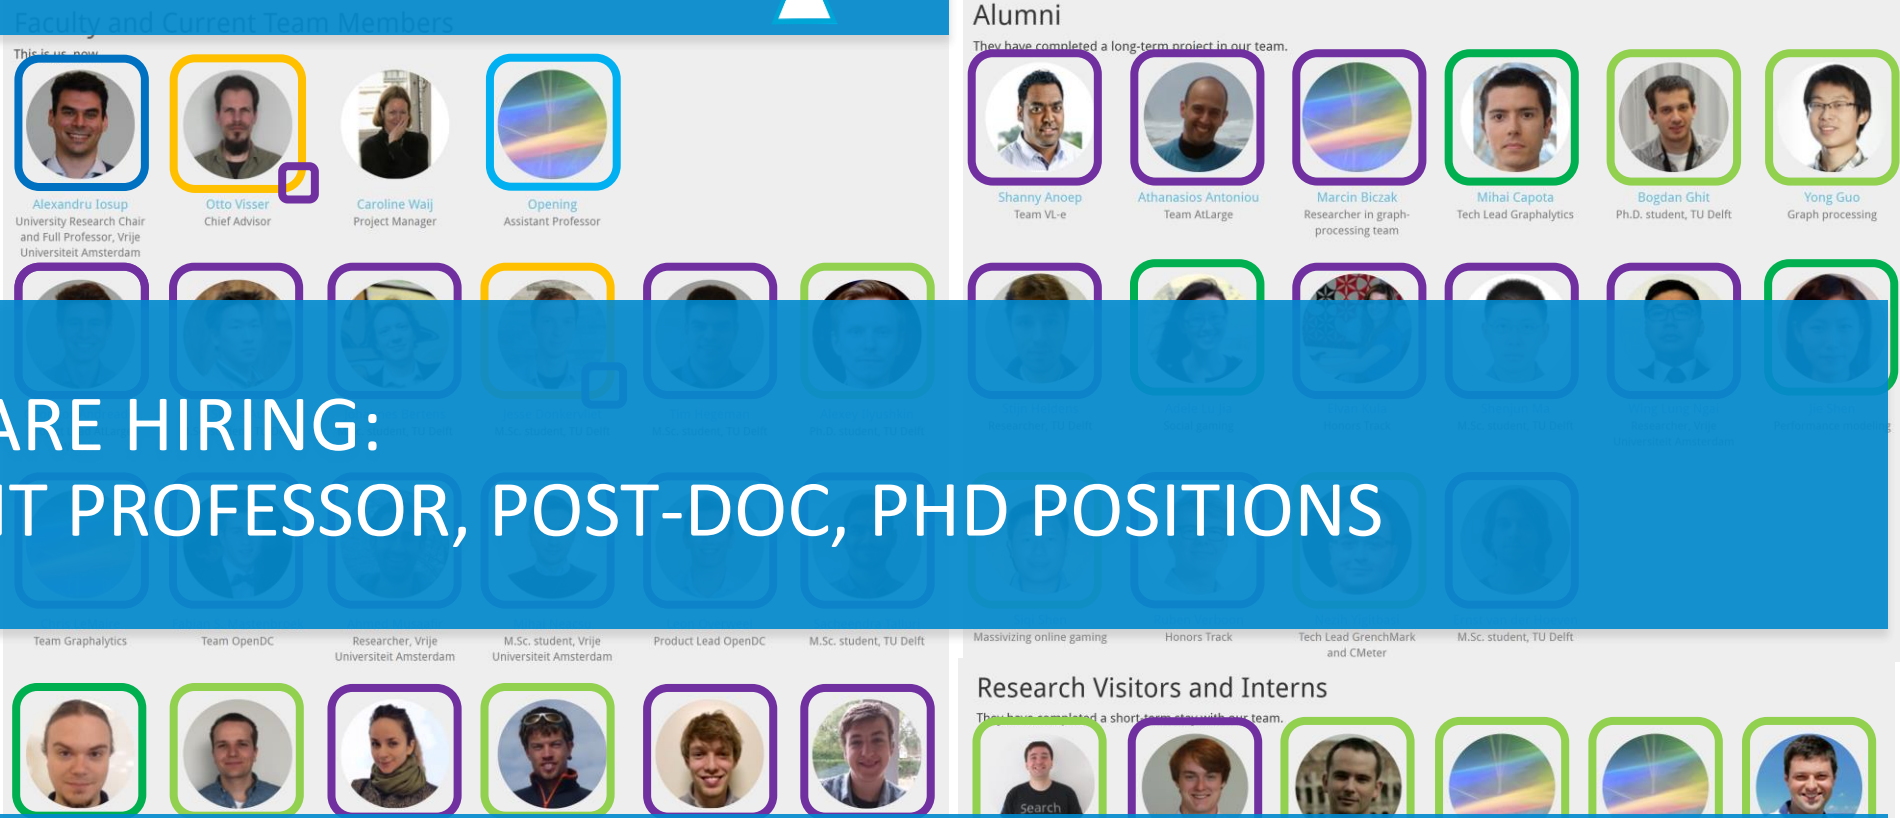
PROF. DR. IR. ALEXANDRU IOSUP

- 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
 - > NL ICT Researcher of the Year
 - > NL Higher-Education Teacher of the Year
 - > NL Royal Young Academy of Arts & Sciences



ATLARGE RESEARCH, OUR TEAM


<http://atlarge.science/people.html>



The screenshot shows a website layout with three main sections: 'Faculty and Current Team Members', 'Alumni', and 'Research Visitors and Interns'. Each section contains a grid of profile pictures. Some profiles are highlighted with colored boxes: a blue box around Alexandru Iosup, a yellow box around Otto Visser, and a purple box around Caroline Wajj. The 'Alumni' section includes profiles for Shanny Anoep, Athanasios Antoniou, Marcin Biczak, Mihai Capota, Bogdan Ghit, and Yong Guo. The 'Research Visitors and Interns' section includes profiles for Search, and several others with rainbow-colored profile pictures.

Professor

Assistant Prof.



WE ARE HIRING:

ASSISTANT PROFESSOR, POST-DOC, PHD POSITIONS

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!

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.





~40'

Massivizing Computer Systems



A Structured Discussion

~3'

About Our Team 




~10'

The Golden Age of Massive Systems ...Yet We Are in a Crisis 

- The main challenges 
- How we address them 

~5'

Massivizing Computer Systems: Examples 

1. The Ecosystem Navigation Challenge 
2. The New World Challenge 
3. The Scheduling Challenge 

~2'

Take-Home Message 

Vote for what
you want to hear
about

VISION:
WHAT DOES OUR SOCIETY NEED?

ISN'T THIS ALREADY
HAPPENING?

prosperous society

blooming economy

inventive academia

wise governance


THIS IS THE GOLDEN AGE OF DISTRIBUTED ECOSYSTEMS




Education for Everyone (Online)




Business Services




Grid Computing



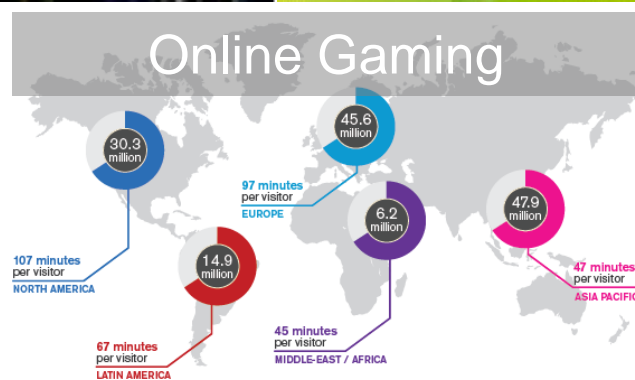
Grid Computing



Big Science



Online Gaming



AVERAGE DAILY ONLINE GAMERS WORLDWIDE

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



Datacenters



Daily Life



THIS IS THE GOLDEN AGE OF DISTRIBUTED COMPUTER SYSTEMS

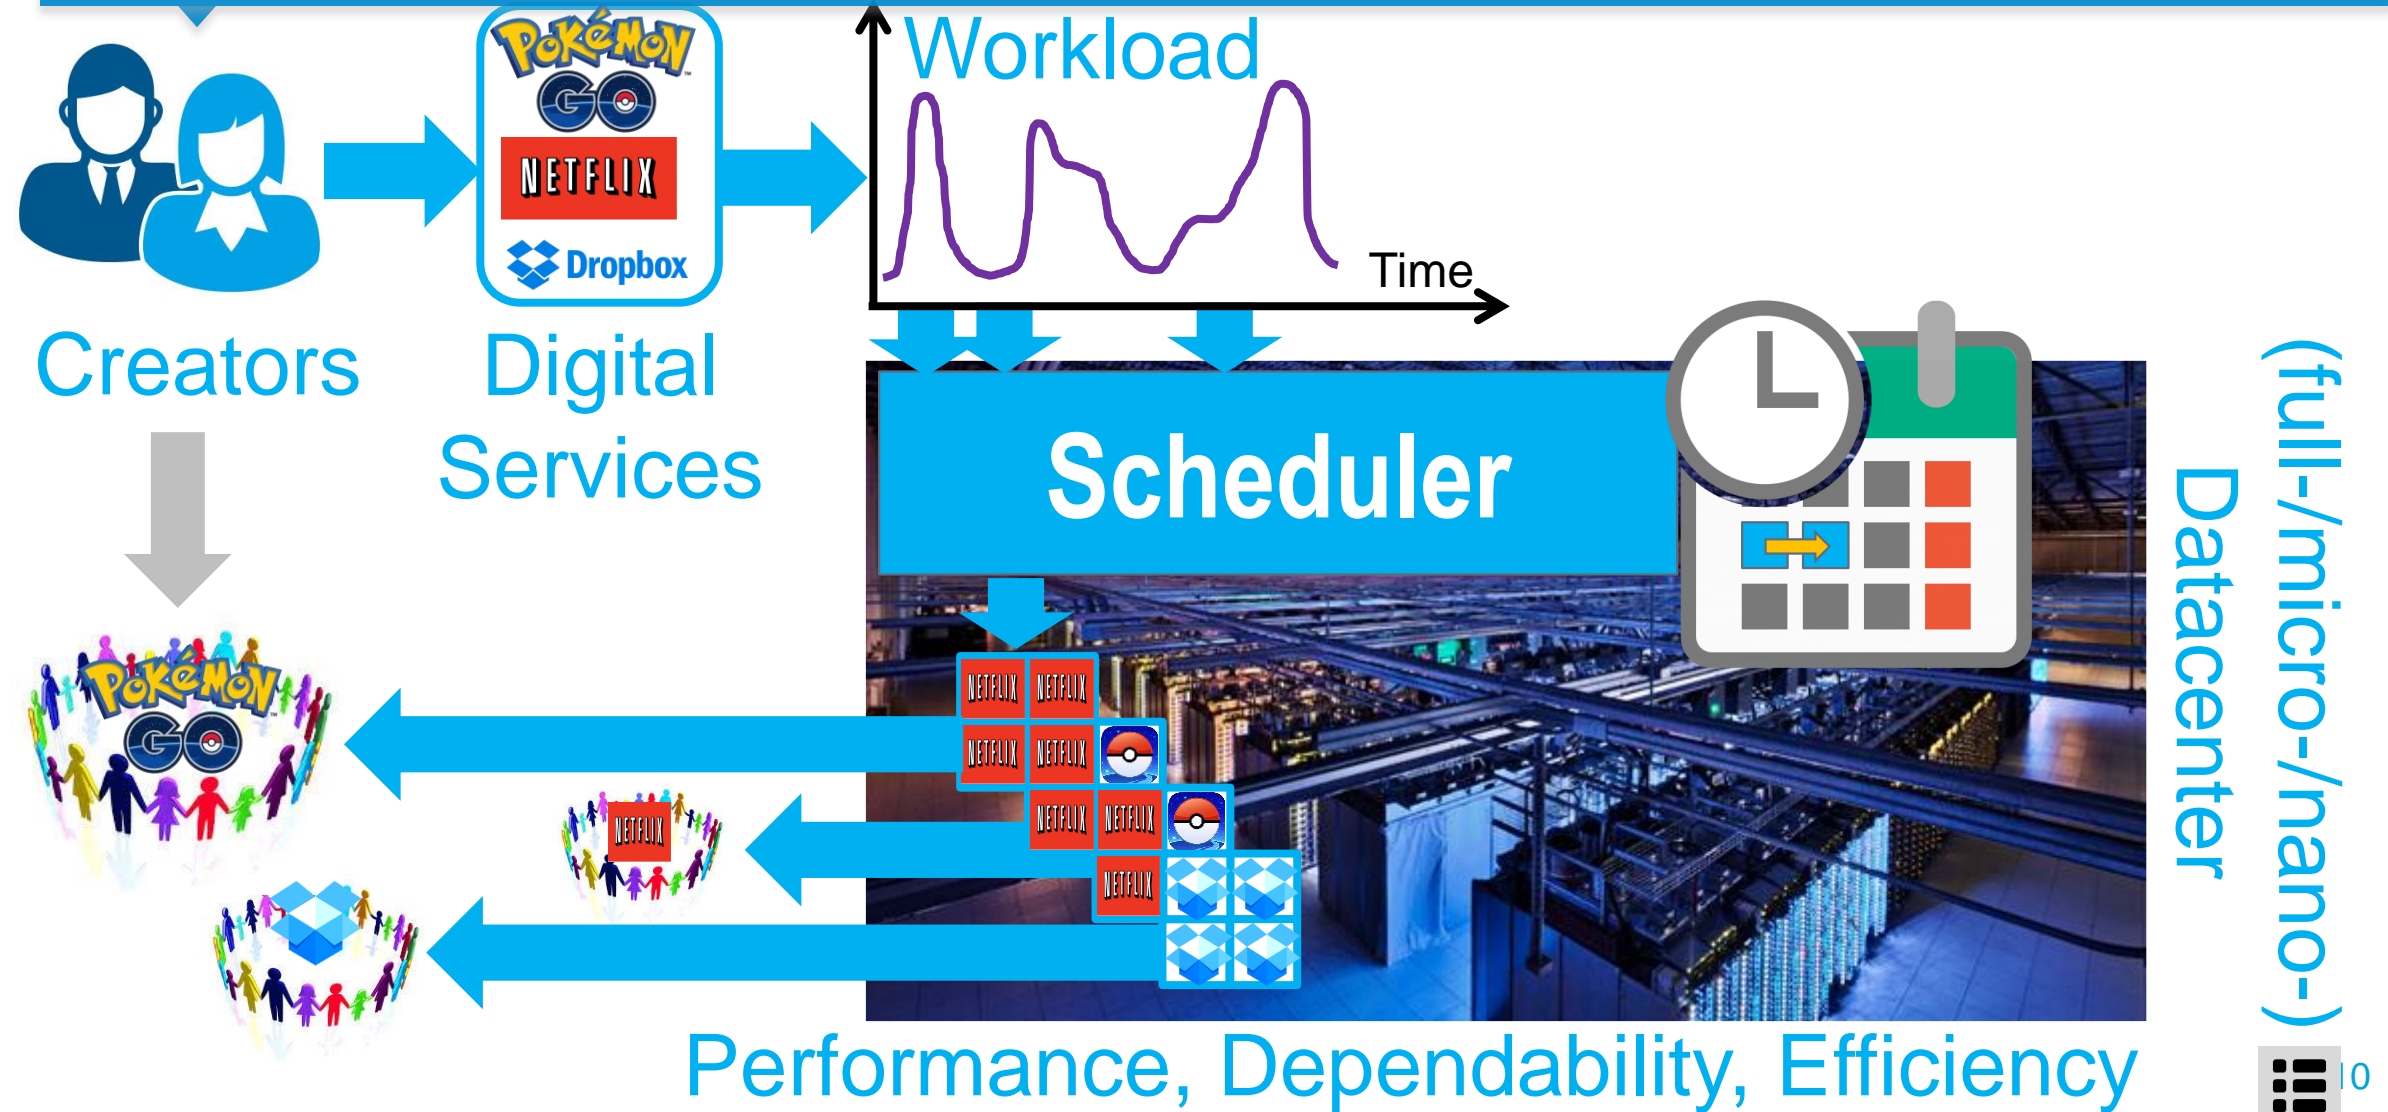
Do you recognize this App?



Here is how it operates...

Daily Life

THE CURRENT TECHNOLOGY STACK: DATACENTER, SCHEDULER



(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.

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
(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 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

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-TRIVIALY FROM ITS SYSTEMS AND USERS (RECURSIVELY)

WHICH APPROACH?

MODERN DISTRIBUTED SYSTEMS AND PROBLEM-SOLVING



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



science + engineering + design





~40'

Massivizing Computer Systems

A Structured Discussion

~3' — About Our Team →

~10' — The Golden Age of Massive Systems ...Yet We Are in a Crisis →

- The main challenges →
- How we address them →

~5' — Massivizing Computer Systems: Examples →

Vote for what
you want to hear
about

1. The Ecosystem Navigation Challenge →
2. The New World Challenge →
3. The Scheduling Challenge →

~2' — Take-Home Message →

MASSIVIZING COMPUTER SYSTEMS

WHICH APPROACH?

MODERN DISTRIBUTED SYSTEMS AND PROBLEM-SOLVING

HOW TAUGHT?

AWARD-WINNING METHOD, BASED ON GAMIFICATION


- Choose your own path:

- > The Ecosystem Navigation Challenge (Understanding + Exp.)
- > The New World Challenge (Abstraction + Design)
- > The Scheduling Challenge (Design + Operation)

The Ecosystem Navigation Challenge

Ecosystem operators: how to prove capabilities? How to tune the tool? In which technology to invest? Which tech to DevOp in-house?

Ecosystem customers: how to choose the right tool? For batch, workflows, stream, transactions, etc. (No one size fits all!)





Alexandru Iosup
Chair

Nikolas Herbst
Vice-Chair

The SPEC RG Cloud Group

Methodology, Benchmarking, and Performance Analysis of Cloud Systems and Applications

“A broad approach, **relevant for both academia and industry**, to cloud benchmarking, quantitative evaluation, and experimental analysis.”

“To develop new **methodological elements** for gaining deeper understanding not only of **cloud performance**, but also of **cloud operation and behavior**”

“... through diverse quantitative evaluation tools”

<http://research.spec.org/working-groups/rg-cloud-working-group.html>

Sales team:
“\$**200/year**
for research
teams”

A General Approach for Ecosystem Benchmarking (Works for Cloud and Big Data Benchmarking)


Q1: What is the performance of production IaaS cloud services?

Q4: What is the performance of graph-processing systems?

Q2: How variable is the performance of widely used production cloud services?

+ many other questions...

Q3: What is the performance of IaaS-cloud auto-scalers?



EXPERIMENTAL METHOD OF DISCOVERY

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

Our Prototypes (*in physico/in vitro*)



+



We also use clouds

+



OpenDC

And simulators (*in silico*)

300+ scientists as users

We won IEEE Scale Challenge 2014

Our Method

Q1

- General performance technique, **adapted to clouds**: model performance of individual components; system performance is performance of workload + model [Saavedra and Smith, ACM TOCS'96]

Iosup et al., Performance Analysis of Cloud Computing Services for Many Tasks Scientific Computing, IEEE TPDS 2011.
Highest cited article in the best journal of the field (2009-2015).

Q2

- Performance traces from CloudStatus, to understand **variability**
 - All Amazon AWS and Google GAE services
 - Periodic performance probes, sampling rate under 2 minutes
- Simulations to assess impact of performance variability
 - Based on traces collected from other applications

Iosup, Yigitbasi, Epema. On the Performance Variability of Production Cloud Services, IEEE CCgrid 2011.
Highest cited study on performance variability in clouds.

Our Method

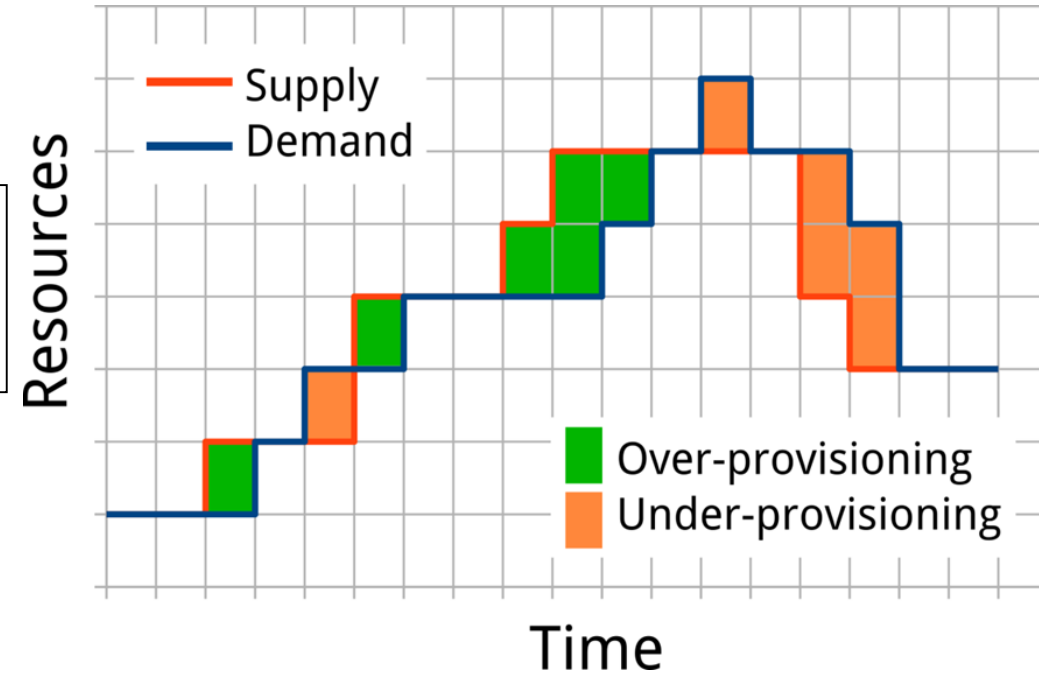
Q3

Ilyushkin et al. An Experimental Performance Evaluation of Autoscaling Policies for Complex workflows. ICPE 2017.
[Nominated for Best Paper Award.](#)

- Real-world experiments (1 workload)
- Later: simulations (more parameters)

Q4

- Many classes of [algorithms](#) used in practice
- Diverse real and synthetic [datasets](#)
 - Recently, Broido and Clauset found that power-laws are rare in graphs
- Diverse set of [experiments](#) representative for practice



Main Findings

Q1

- Lower performance than theoretical peak in IaaS services
 - Especially CPU (GFLOPS)
 - Not explained by traditional models, not covered by SLA

Q2

- Performance variability in IaaS and PaaS services
 - Explored in longitudinal study of Amazon Web Services and Google App Engine
 - Not captured in traditional models

- Compared performance of IaaS clouds with many commercial alternatives, such as supercomputers and clusters

Main Findings

Q3

- Explored impact of auto-scaler (+ ecosystem of schedulers) on over 10 facets of elasticity
 - Findings not explained by traditional models

Q4

- The HPAD model for the performance of graph-processing systems ~ replaces previous theories
 - Performance is function of Platform, but also Dataset and Algorithm
 - With configurable many-/multi-cores, the Hardware also is also crucial for performance

MASSIVIZING COMPUTER SYSTEMS

WHICH APPROACH?

MODERN DISTRIBUTED SYSTEMS AND PROBLEM-SOLVING

HOW TAUGHT?

AWARD-WINNING METHOD, BASED ON GAMIFICATION



- Choose your own path:

- > The Ecosystem Navigation Challenge (Understanding + Exp.)
- > The New World Challenge (Abstraction + Design)
- > The Scheduling Challenge (Design + Operation)

The New World Challenge




Ecosystem operators: new value-adding services, new workloads including FaaS, DevOps workloads



CTAAAGATGATCTTAGTCCCGGTTGAA
TCTTTAGTCCCGGTTGATAACCAACC
GTAATACCAACCGGACTAAAGATCCGG
GGGACTAAAGTCCACCCCTATATATG

TTCAAAATTTCTCAAAAAAGAGGGGAG
GTGATTACATACAATOGGAGGTSCCTA
TTTGTACATACTACATTTGCACCTATGTTT
GTAAGTTGATGAGA



Need Operational Models

Ecosystem customers: new apps, new services, micro-services, customers can become operators (recursive value-chain)





Erwin
van Eyk

Alexandru
Iosup



Serverless / FaaS Execution

Vision and Architecture for Serverless Execution in Cloud Environments

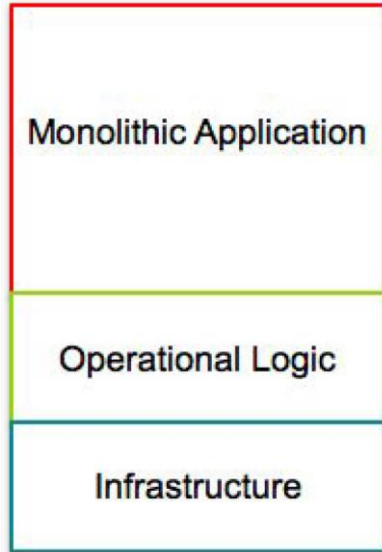


van Eyk, Toader, Talluri, Versluis, Uta, Iosup: Serverless is More: From PaaS to Present Cloud Computing. IEEE Internet Computing Sep/Oct 2018. [[online](#)]

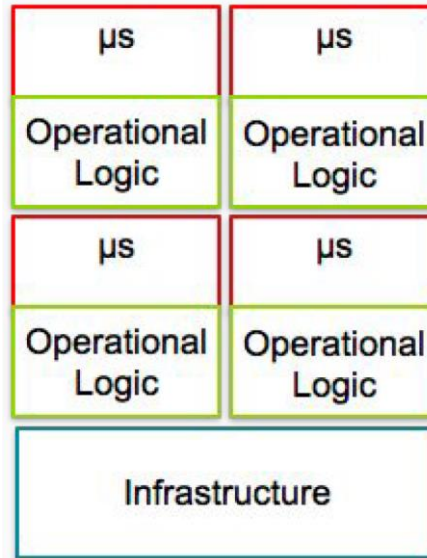
Erwin Van Eyk, Alexandru Iosup, Cristina L. Abad, Johannes Grohmann, Simon Eismann: A SPEC RG Cloud Group's vision on the Performance Challenges of FaaS Cloud Architectures. ICPE 2018. [[online](#)]

Erwin van Eyk, Simon Seif (SAP), Markus Thoenes (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](#)]

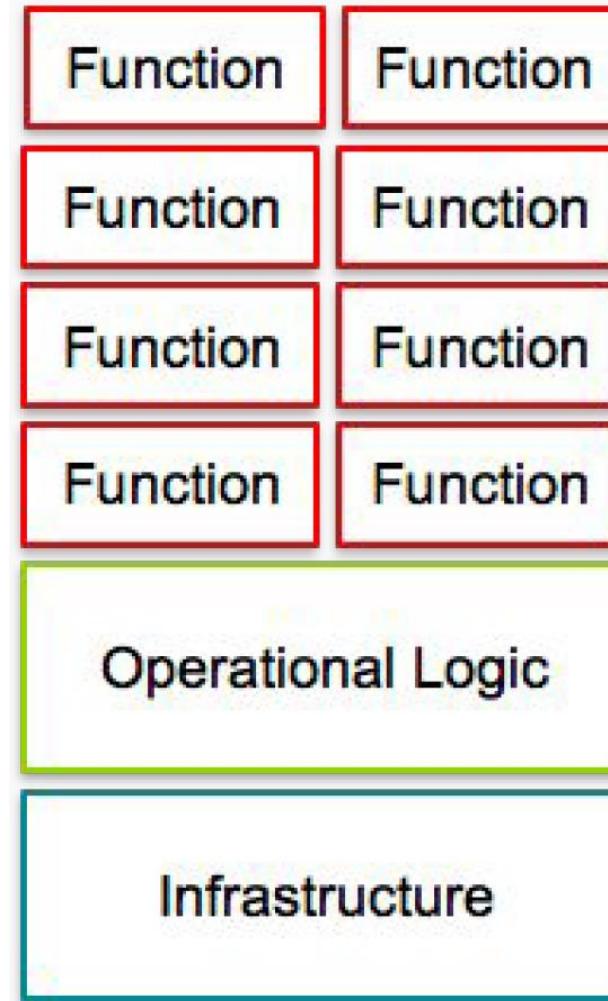
From Monoliths to Microservices to FaaS



- Difficult to Scale
- Infrequent
- Inflexible
- Complex deployment
- Tightly coupled stack

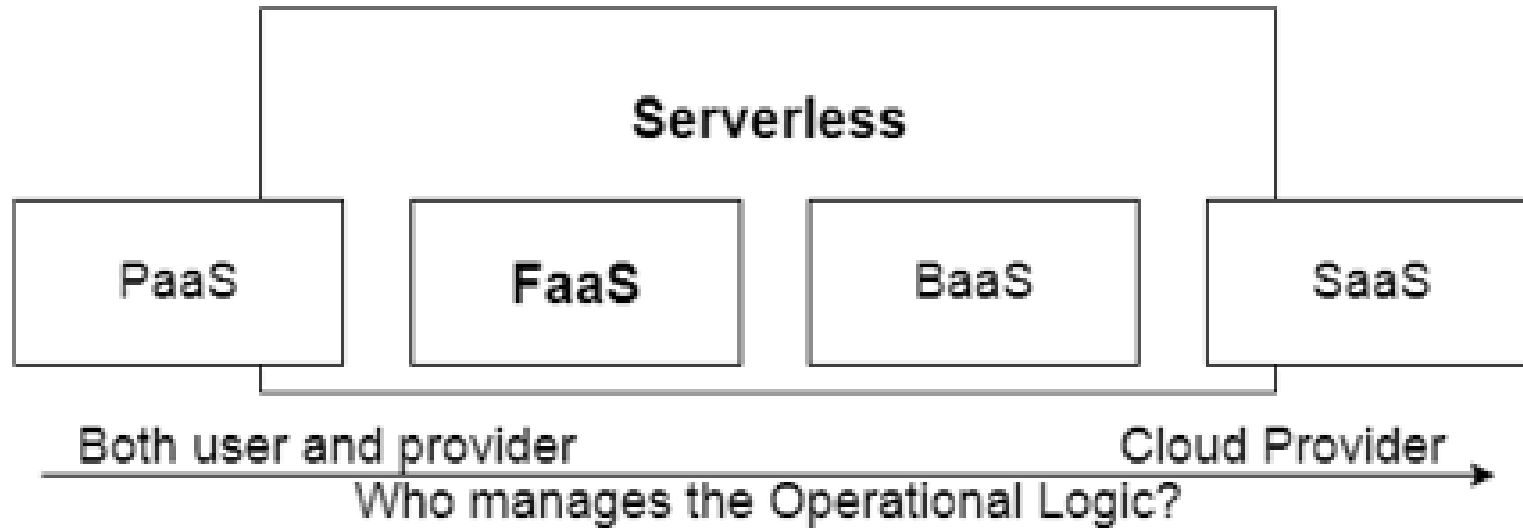


- Scalable
- Frequent
- Flexible
- Complexity: from application logic to operational logic.
- Need for DevOps



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

Serverless and FaaS



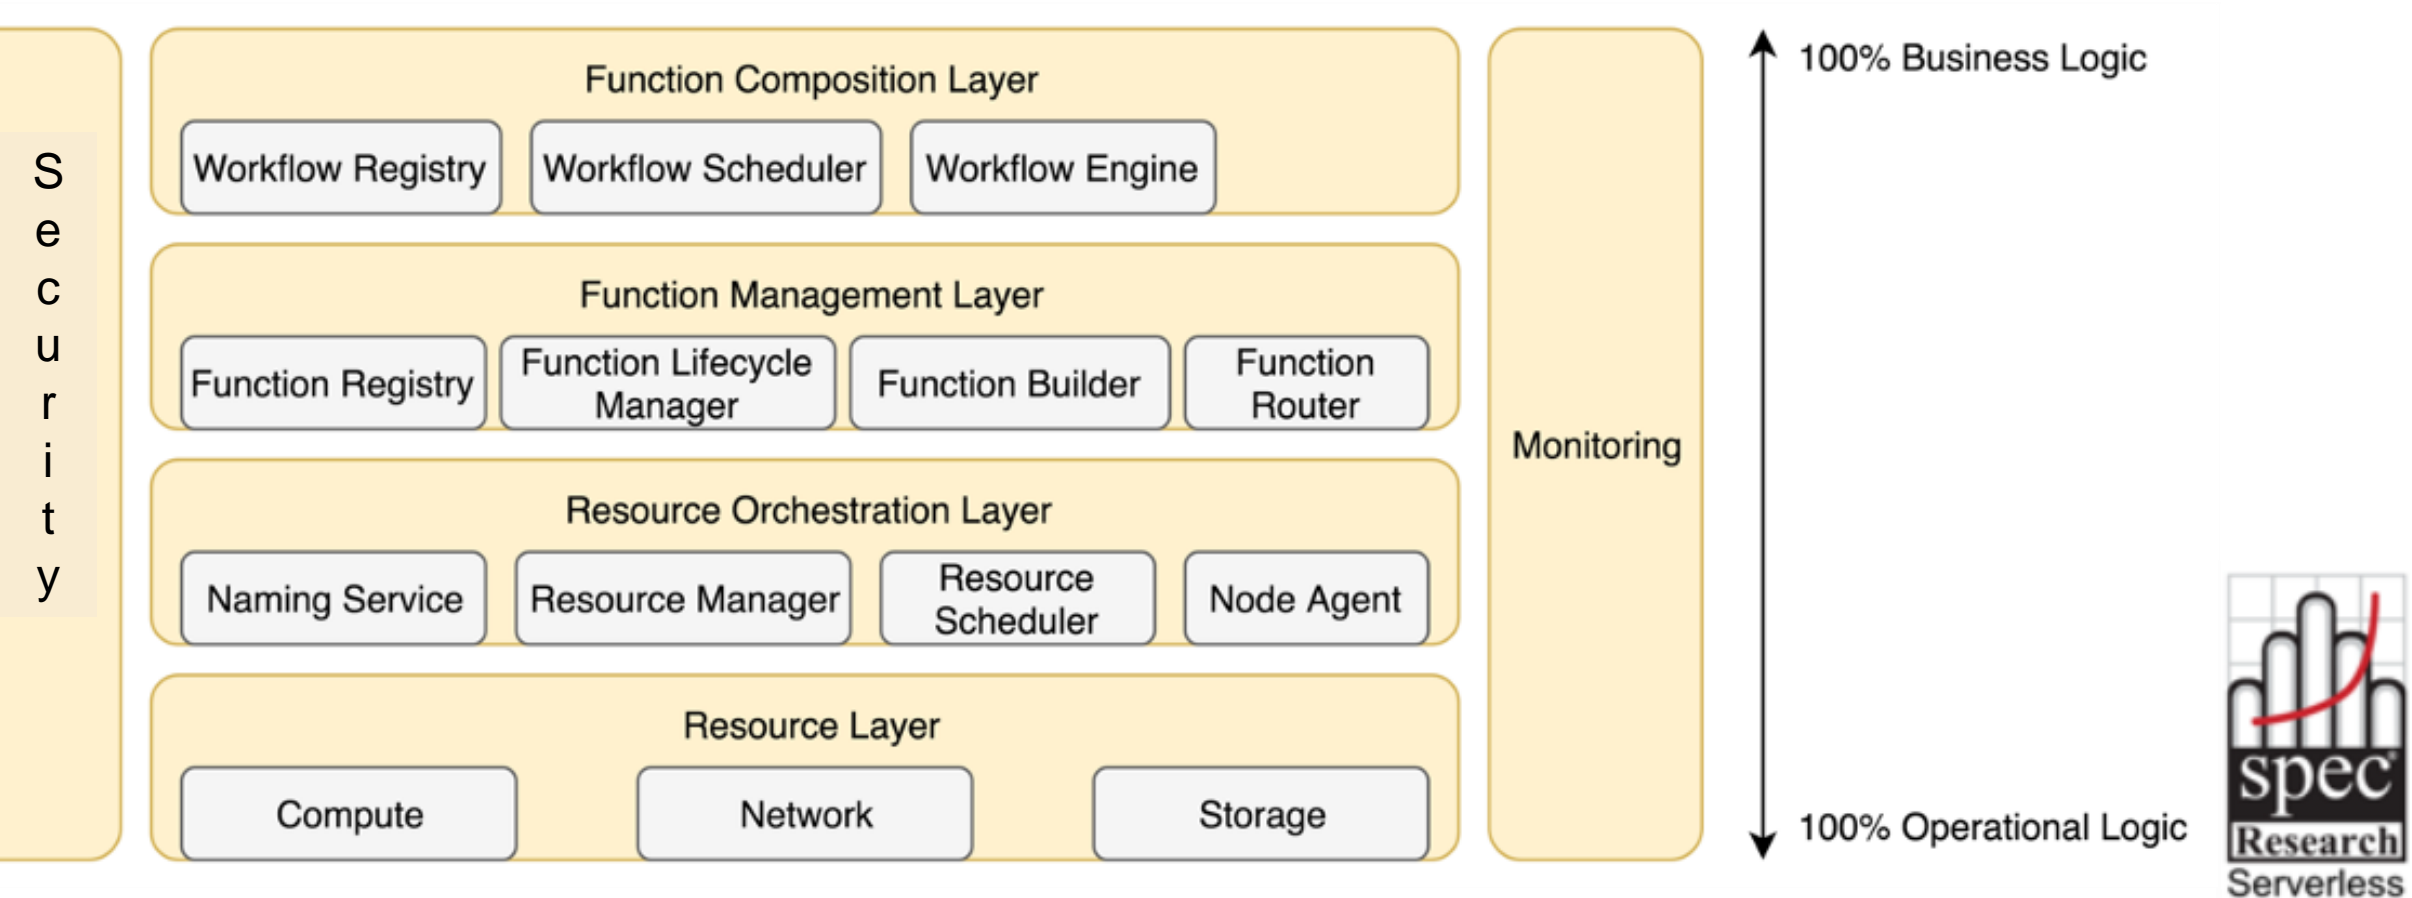
- **Serverless**

- (Almost) no operational logic
- Event-Driven
- Granular billing

- **FaaS**

- A form of *serverless computing*
- User provides a function
- Function deployed and managed by cloud provider

Reference Architecture for FaaS Management



Workflow Management Architecture in Fission.io

Designed by Erwin van Eyk during internship at Platform9, in collaboration w/ Platform9 team and Alexandru Iosup.

1 Core Function / 2 API Server

- Exposes all actions through API

3 Event Store / 4 Projector


- Events update the workflow
- Store has Pub/Sub functionality
- Projector builds current state

5 Fission Proxy

- API access to Fission FaaS

6 Controller / 7 Scheduler

- Workflow manager



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

MASSIVIZING COMPUTER SYSTEMS

WHICH APPROACH?

MODERN DISTRIBUTED SYSTEMS AND PROBLEM-SOLVING

HOW TAUGHT?

AWARD-WINNING METHOD, BASED ON GAMIFICATION

- Choose your own path:

- > The Ecosystem Navigation Challenge (Understanding + Exp.)
- > The New World Challenge (Abstraction + Design)
- > The Scheduling Challenge (Design + Operation)

The Scheduling Challenge



Ecosystem operators:

**Which resources to lend?
Where to place? Penalty v reward?**

**Need scheduling policies for both
users and operators**

Ecosystem customers:

**Which resources to lease?
When? How many? When stop?
Utility functions?**





Vincent
van Beek




Tim
Hegeman



Jesse
Donkervliet



Alexandru
Iosup



Portfolio Scheduling for DCs


Self-Expressive Management of Business-Critical Workloads in Virtualized Datacenters

van Beek, Donkervliet, Hegeman, Hugtenburg, Iosup. Self-Expressive Management of Business-Critical workloads in virtualized Datacenters. IEEE Computer 48(7): 46-54 (2015)

Deng, Song, Ren, Iosup. Exploring portfolio scheduling for long-term execution of scientific workloads in IaaS clouds. SC 2013: 55:1-55:12

Portfolio Scheduling, In A Nutshell

- Datacenters cannot work without one or even several schedulers
- Instead of ephemeral, risky schedulers, we propose to



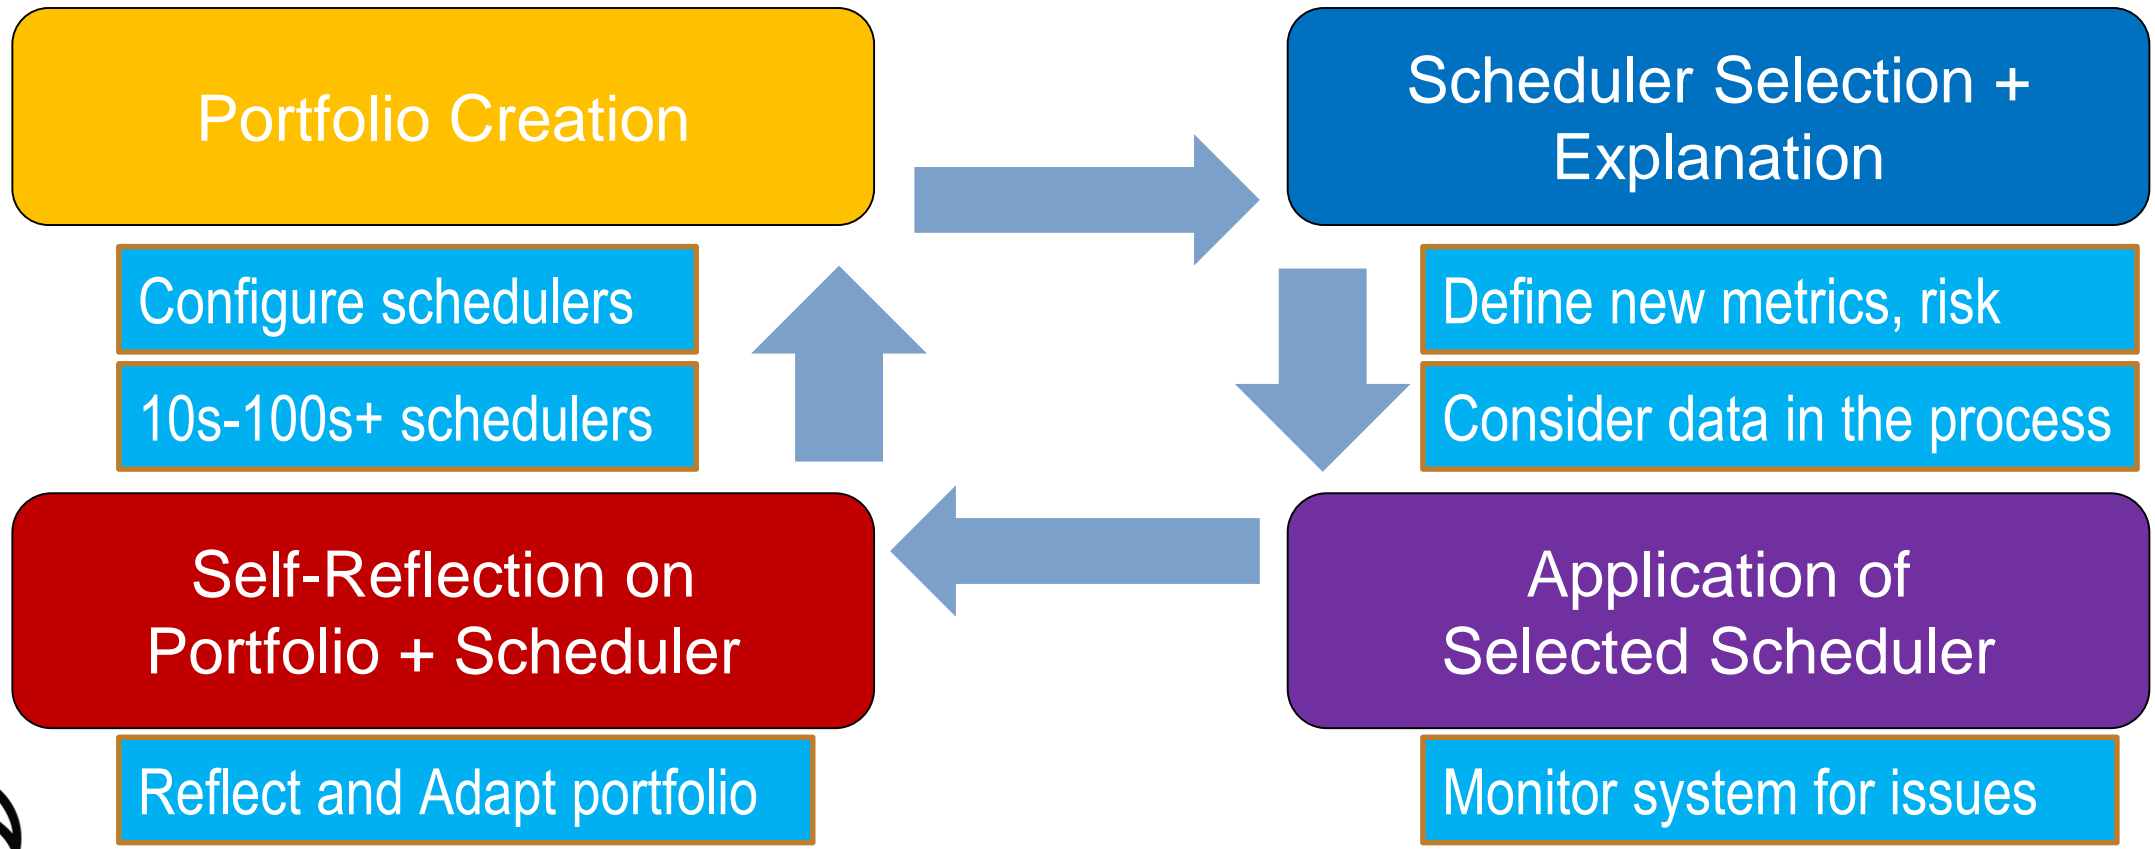
1. Create a set of schedulers (resource provisioning and allocation policies)
2. Select active scheduler online, apply for the next period, analyze results

(Repeat)

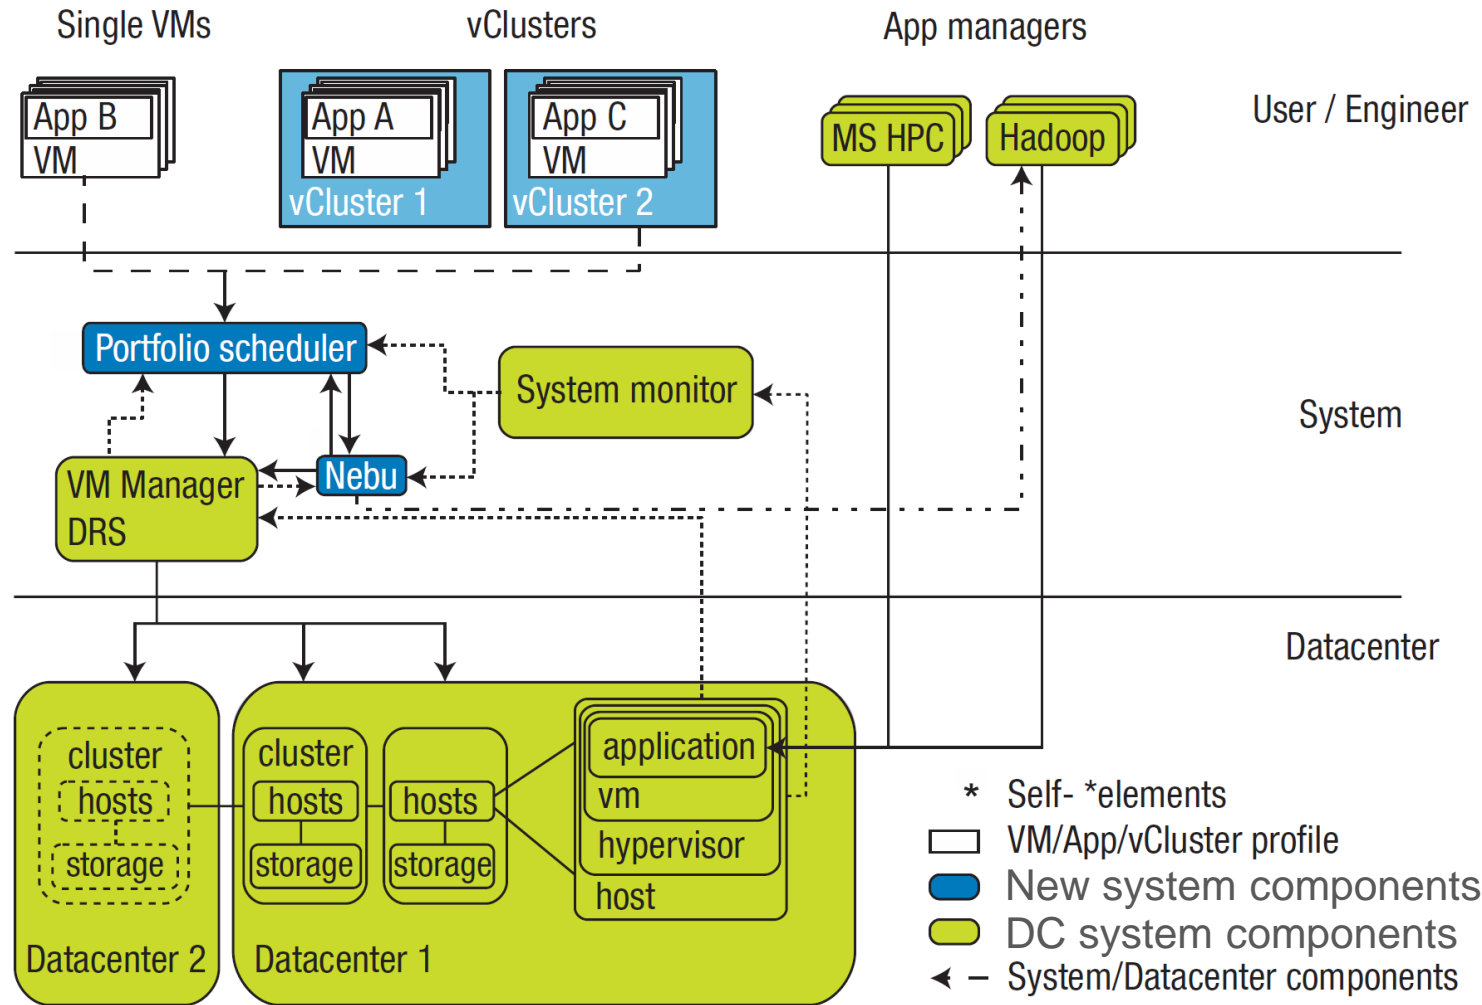
K. Deng et al. Exploring portfolio scheduling for long-term execution of scientific workloads in IaaS clouds. SC'13

Portfolio Scheduling for Computer Systems

Portfolio Scheduling



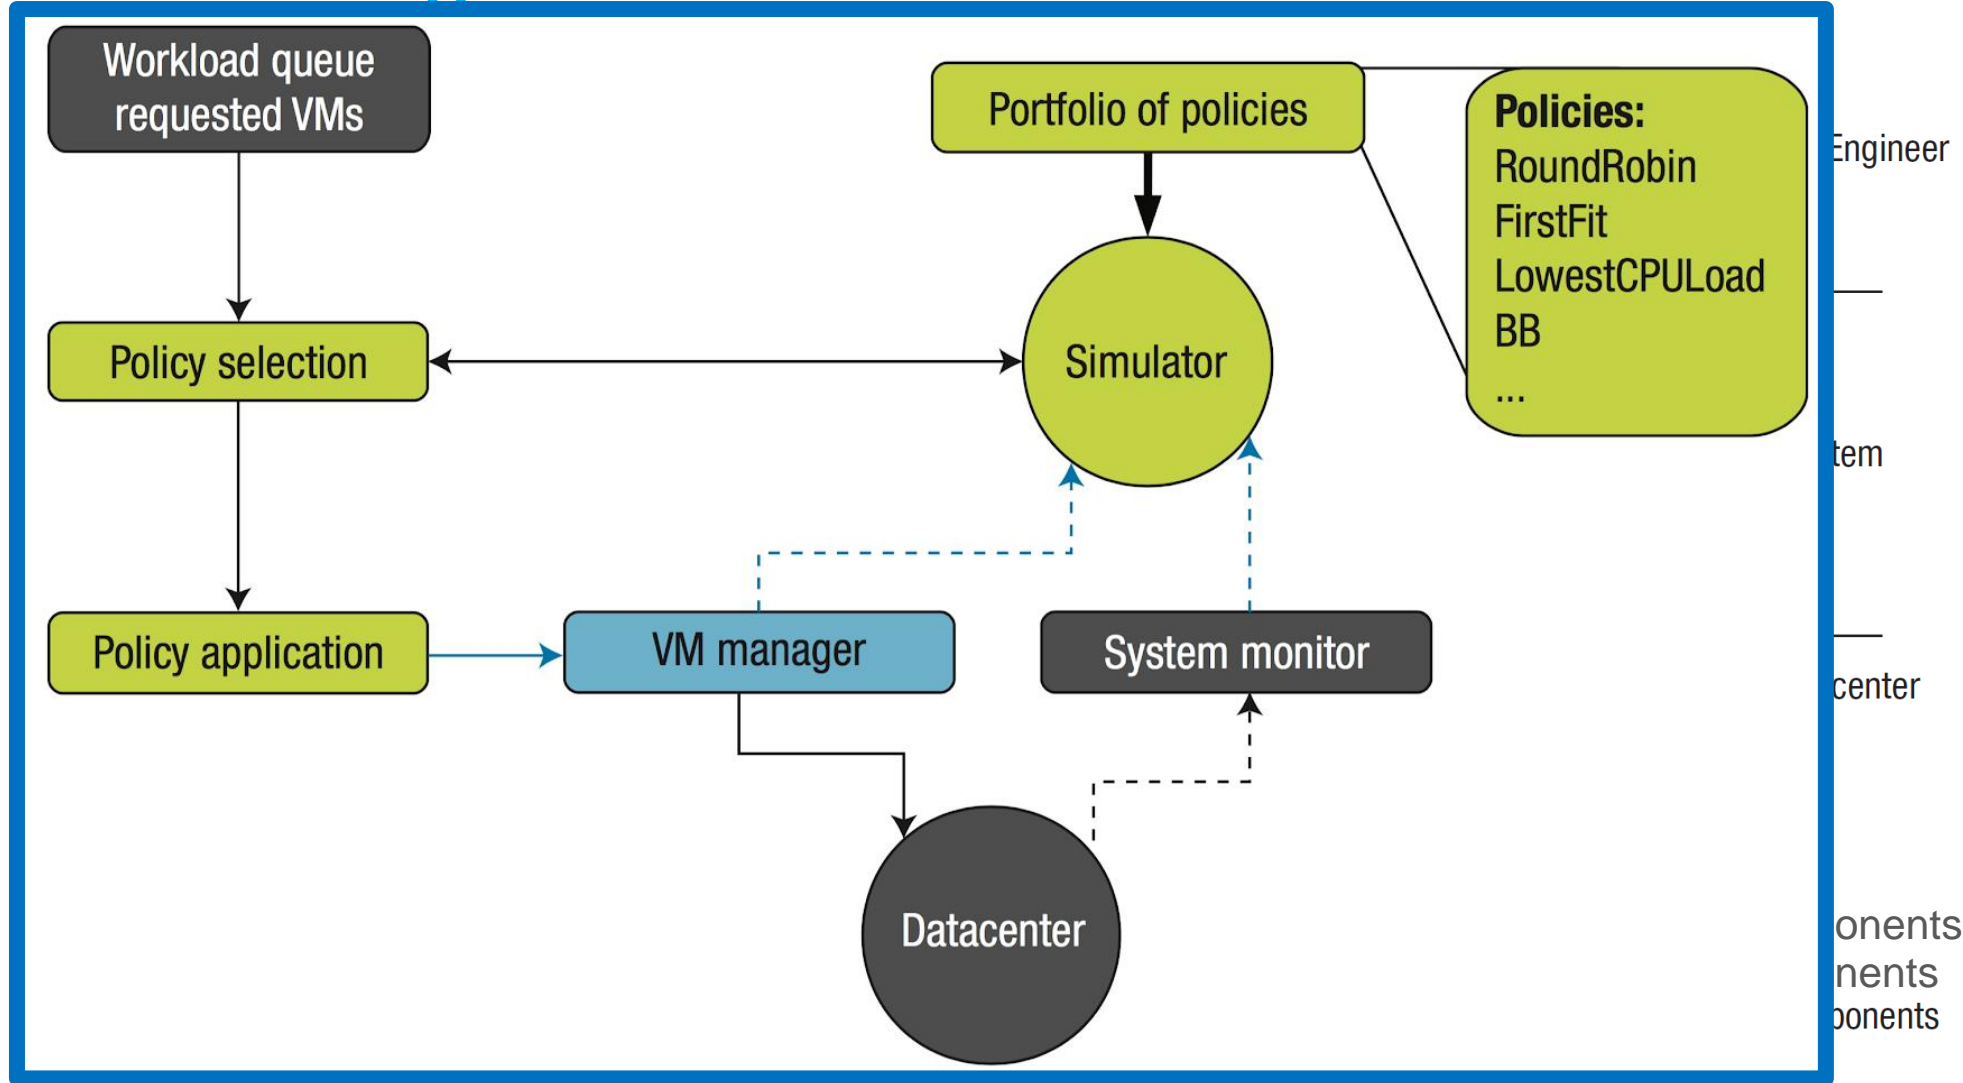
Portfolio Scheduling in Datacenters



V. van Beek et al. Mnemos: Self-Expressive Management of Business-critical workloads in Virtualized Datacenters. IEEE Computer 2015



Portfolio Scheduling in Datacenters

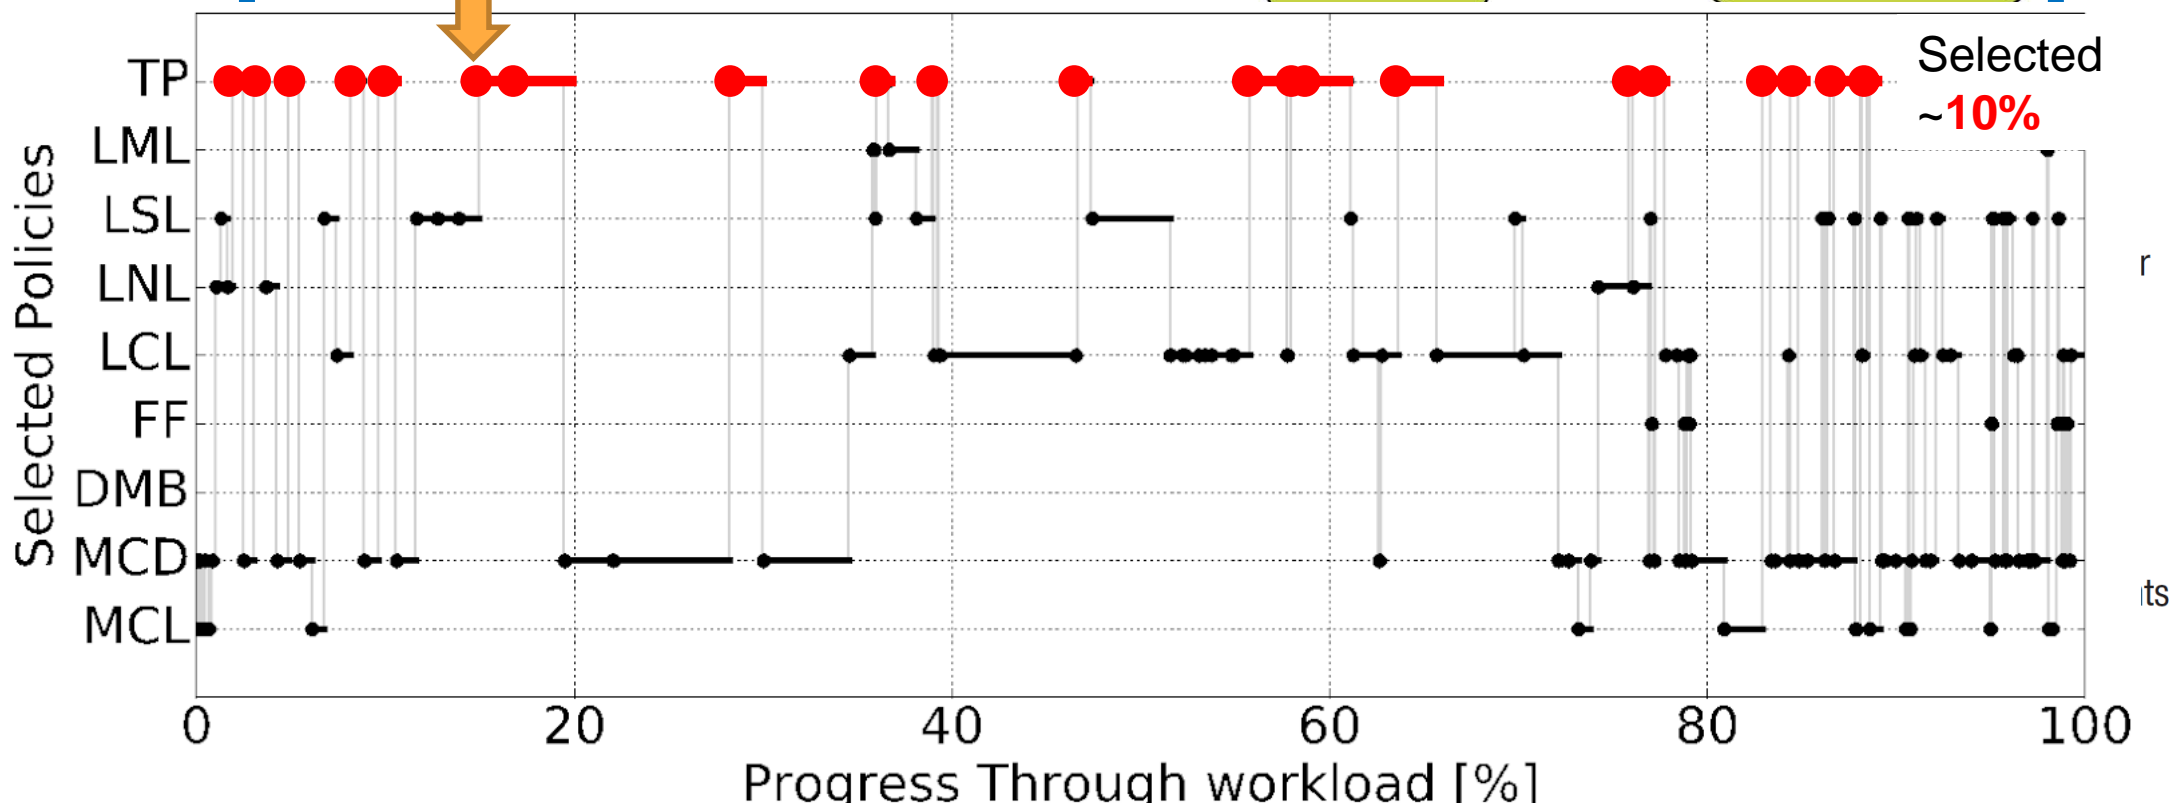
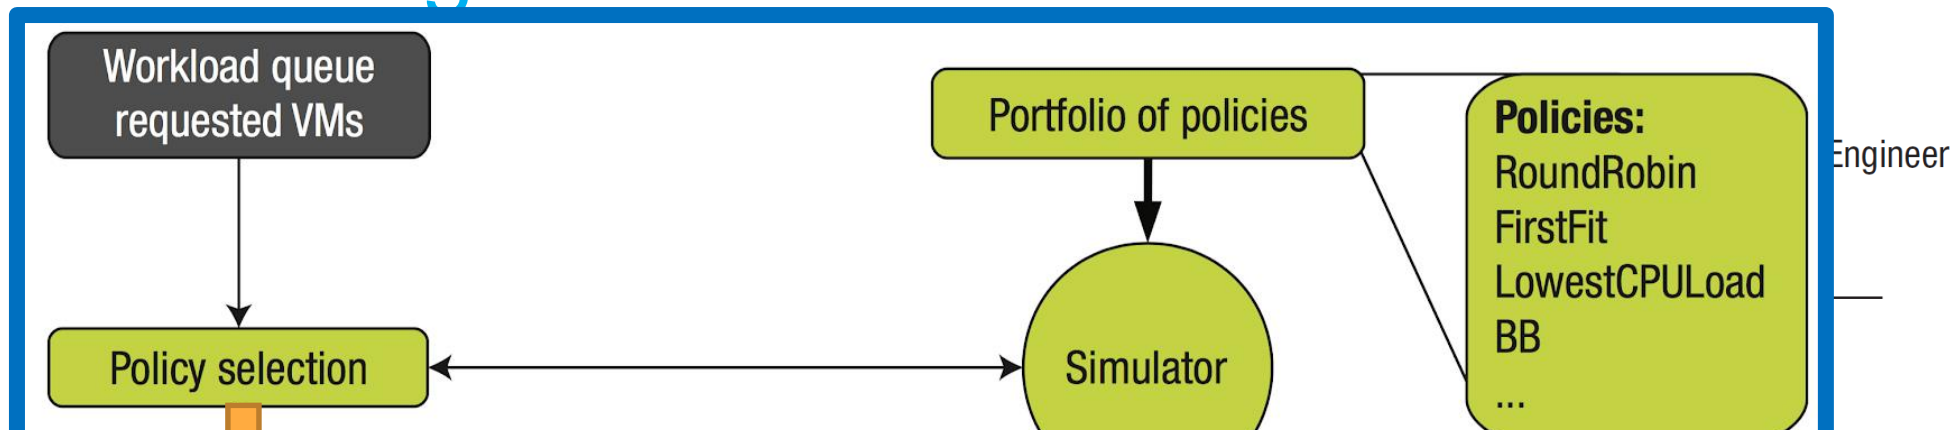


V. van Beek et al. Mnemos: Self-Expressive Management of Business-critical workloads in virtualized Datacenters. IEEE Computer 2015


←···· API Commands
 ←···· Monitoring data
 ←···· Nebu application connection



Portfolio Scheduling in Datacenters



Portfolio Scheduling in Datacenters



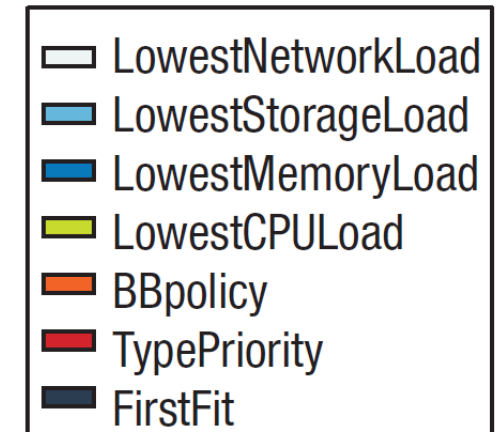

1. Portfolio scheduler achieves the lowest risk of all scenarios.
2. Portfolio scheduler achieves at least **35%** lower DRR compared to individual policies.
3. **40%** lower DRR than commercial production system (REPLAY).

Portfolio Scheduling in Datacenters

Not performance-related, but: A portfolio scheduler can explain each decision by presenting its decision data.

Q: Can our sysadmin do this? Can we? (Rhetorical)

Policy selected,
fraction of decisions



Scenario DC

Unknown: Fundamental limitations? Use in ecosystem? User studies?

V. van Beek et al. Mnemos: Self-Expressive Management of Business-critical workloads in Virtualized Datacenters. IEEE Computer 2015





~40'

Massivizing Computer Systems




A Structured Discussion

~3' — About Our Team 

~10' — The Golden Age of Massive Systems ...Yet We Are in a Crisis 

- The main challenges 
- How we address them 

~5' — Massivizing Computer Systems: Examples 

1. The Ecosystem Navigation Challenge 
2. The New World Challenge 
3. The Scheduling Challenge 

Vote for what
you want to hear
about

~2' — Take-Home Message 

MASSIVIZING COMPUTER SYSTEMS

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



- 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
- Much left to do, as we are merely beginning ...
 - You can help! You can make a career, in science / industry.


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.losup@vu.nl 

+31-20 59 89468 (Amsterdam) 

@Alosup 

<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

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:**
 - > Understand and Create
 - > Experiment and Operate
- **ICT = ecosystems:**
 - > Utilities and services
 - > Automated
 - > Efficient
 - > Controlled
 - > Ecosystems
 - > Human-guided

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

THE COMPLEXITY CHALLENGE

1. 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)

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



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

THE RESOURCE MANAGEMENT CHALLENGE

Based on Jav 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



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



Radio

Microwave

Infrared

Visible

Ultraviolet

X-Ray

Gamma Ray

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

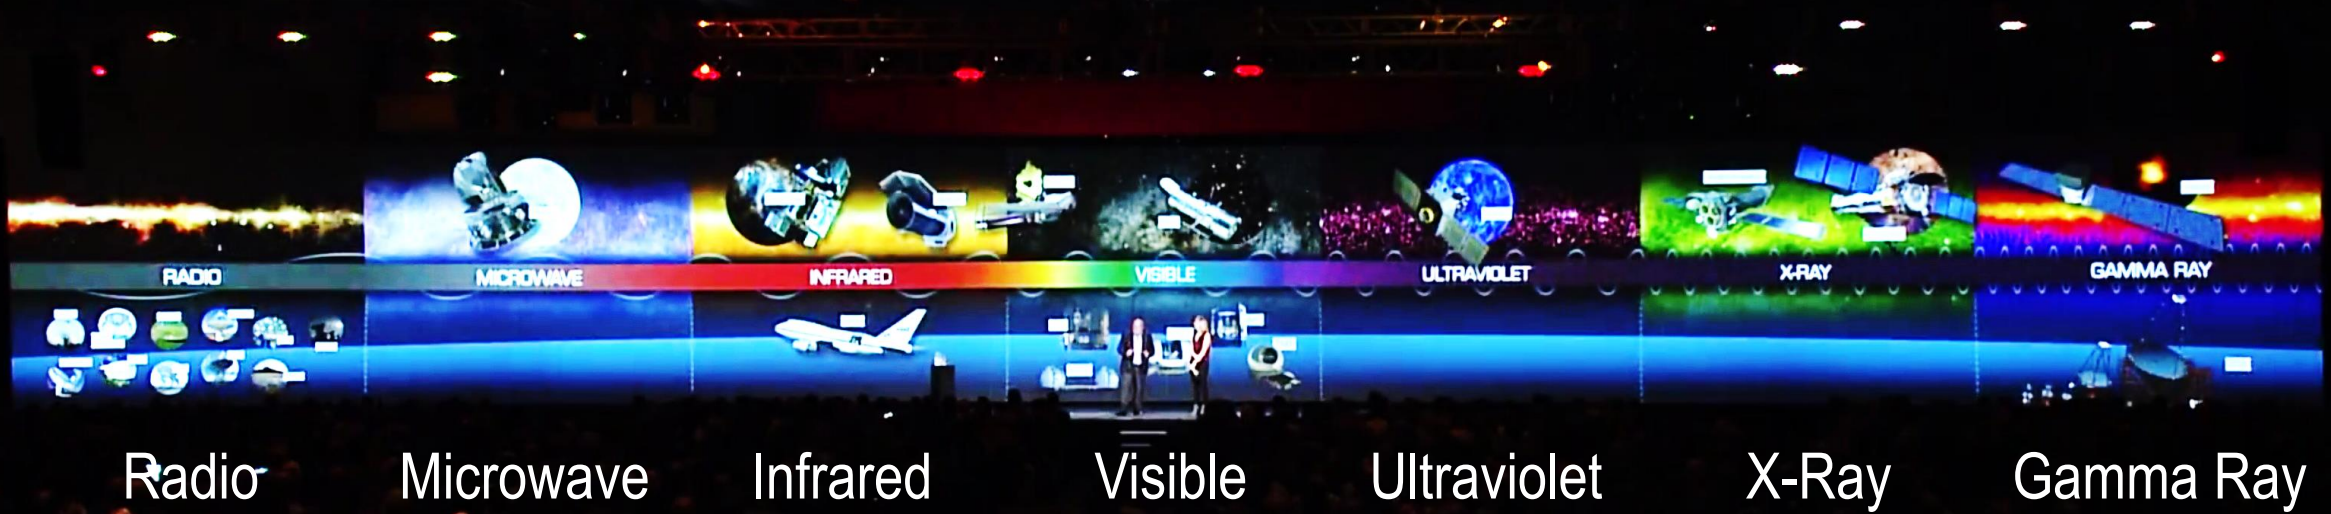
The Square Kilometer Array Factsheet, How much will it cost?, 2012 [[online](#)]

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]



[Ghiț et al. CCGRID'14]



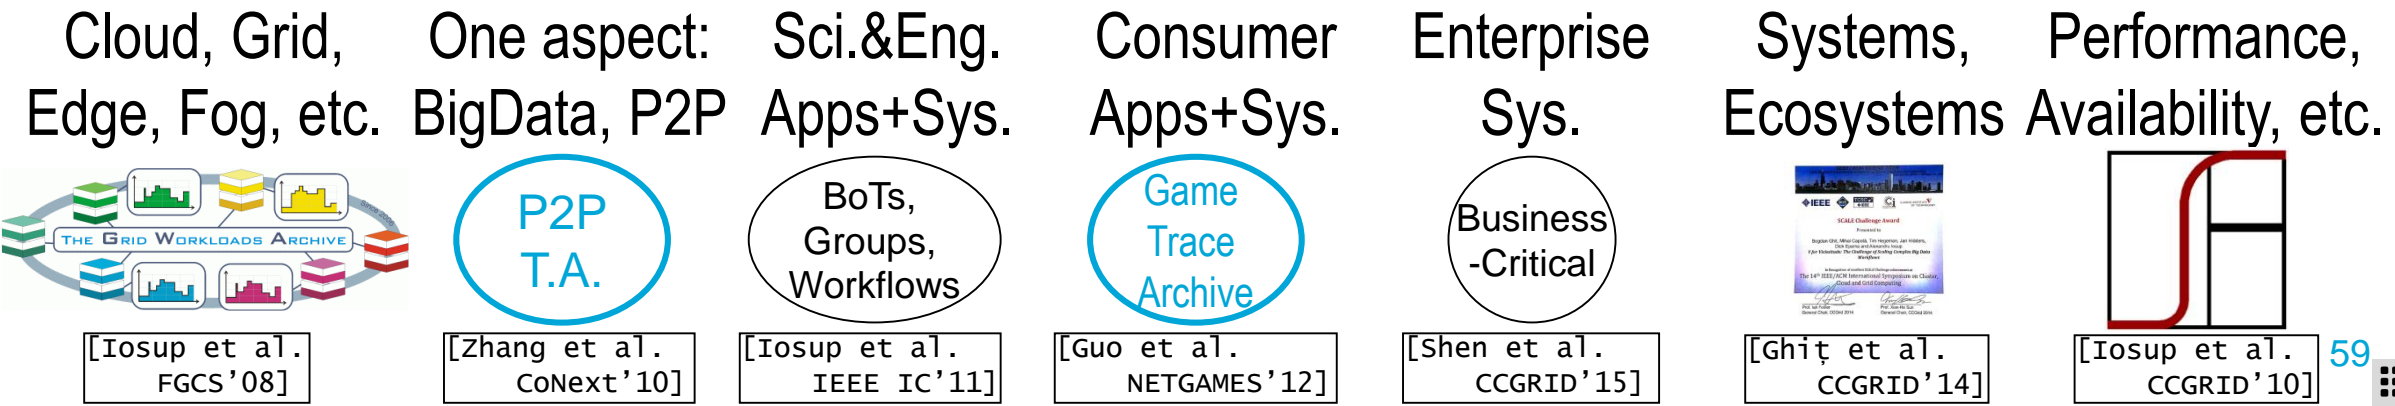
[Iosup et al. CCGRID'10]



MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

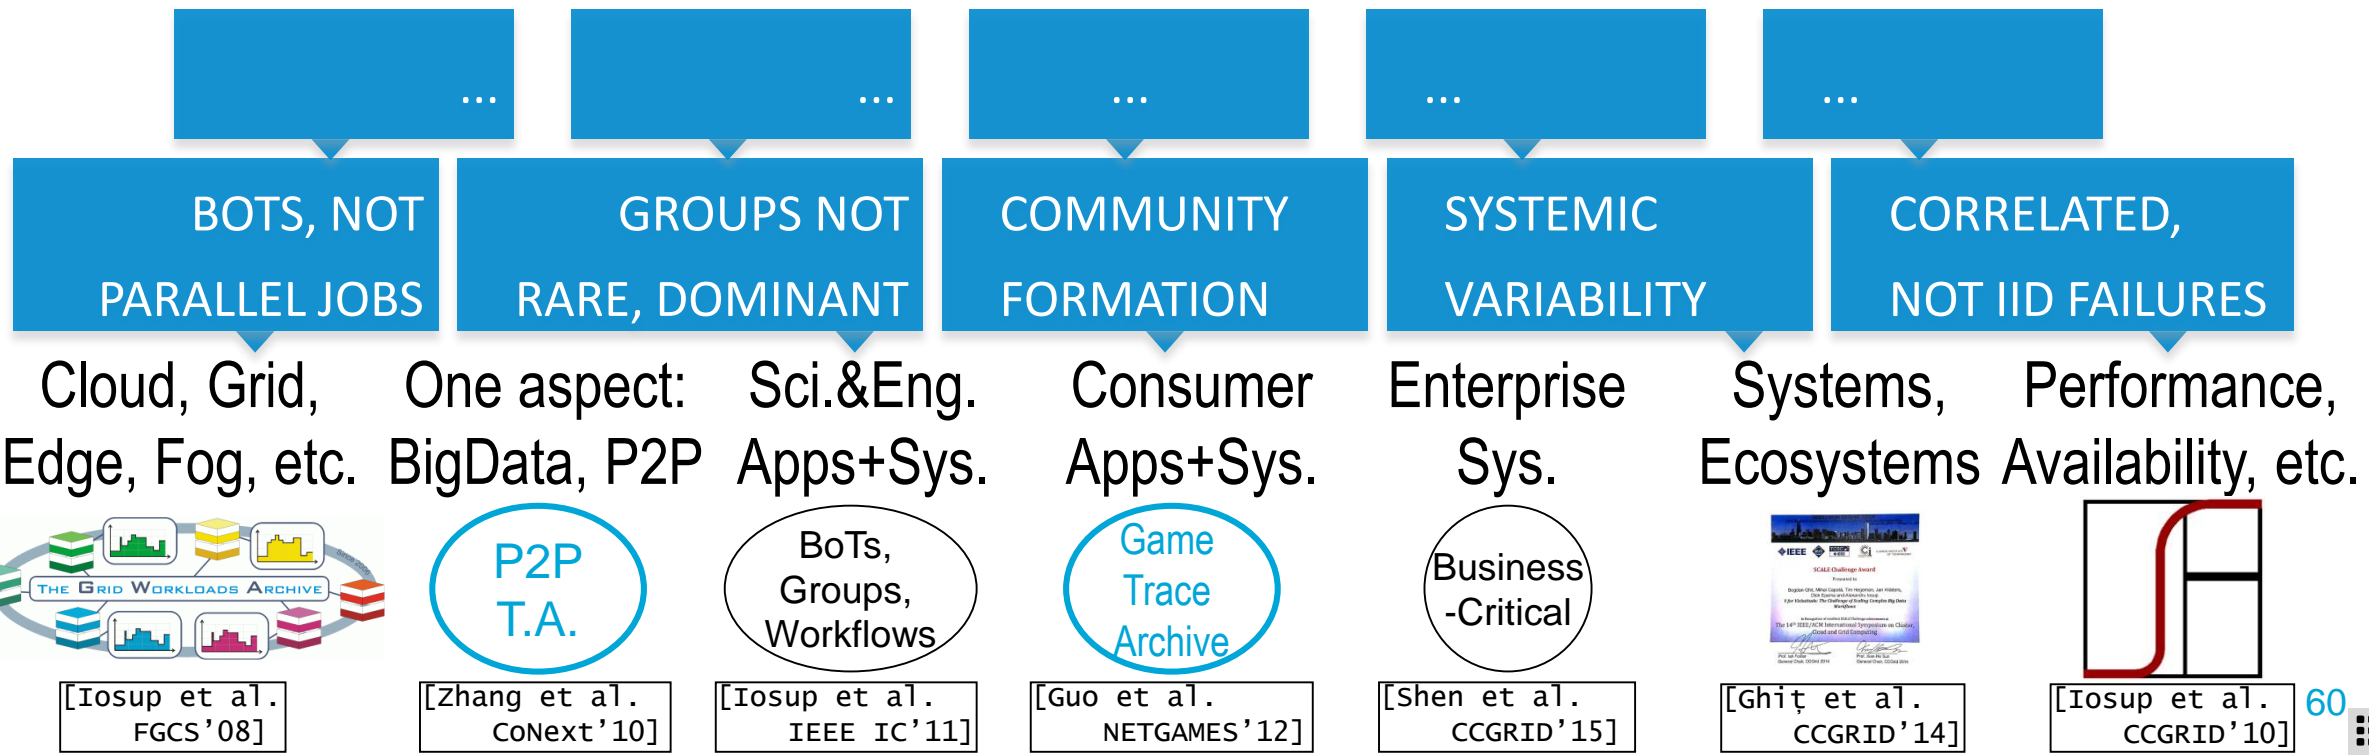
BUT ... WHY WOULD YOU NEED TO UNCOVER AN ARTIFICIAL UNIVERSE?! YOU BUILT IT!



MEANINGFUL DISCOVERY

UNCOVERING THE MYSTERIES OF OUR UNIVERSE, PHYSICAL AND DIGITAL

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



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]



[Ghiț et al.
CCGRID'14]



[Iosup et al.
CCGRID'10]



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. [\[Online\]](#)

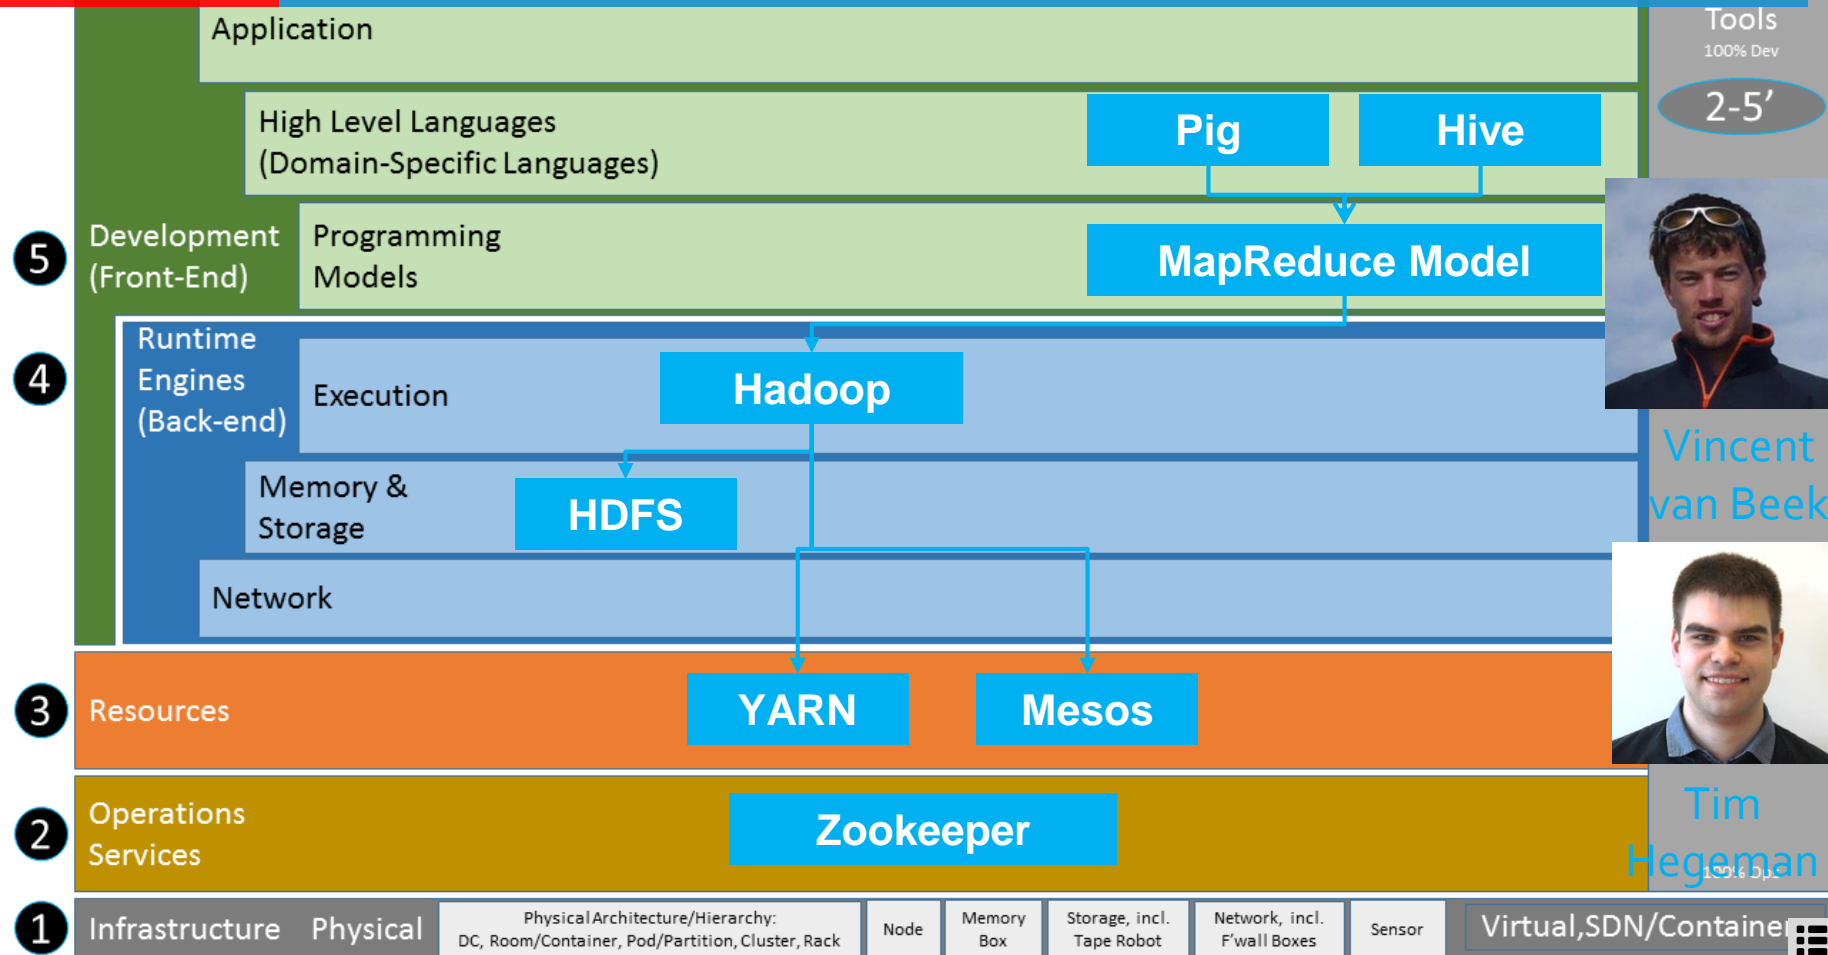
MEANINGFUL DISCOVERY

THE COMPLEXITY CHALLENGE

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS

Focus on Applications,
5 Core Layers:

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



MEANINGFUL DISCOVERY

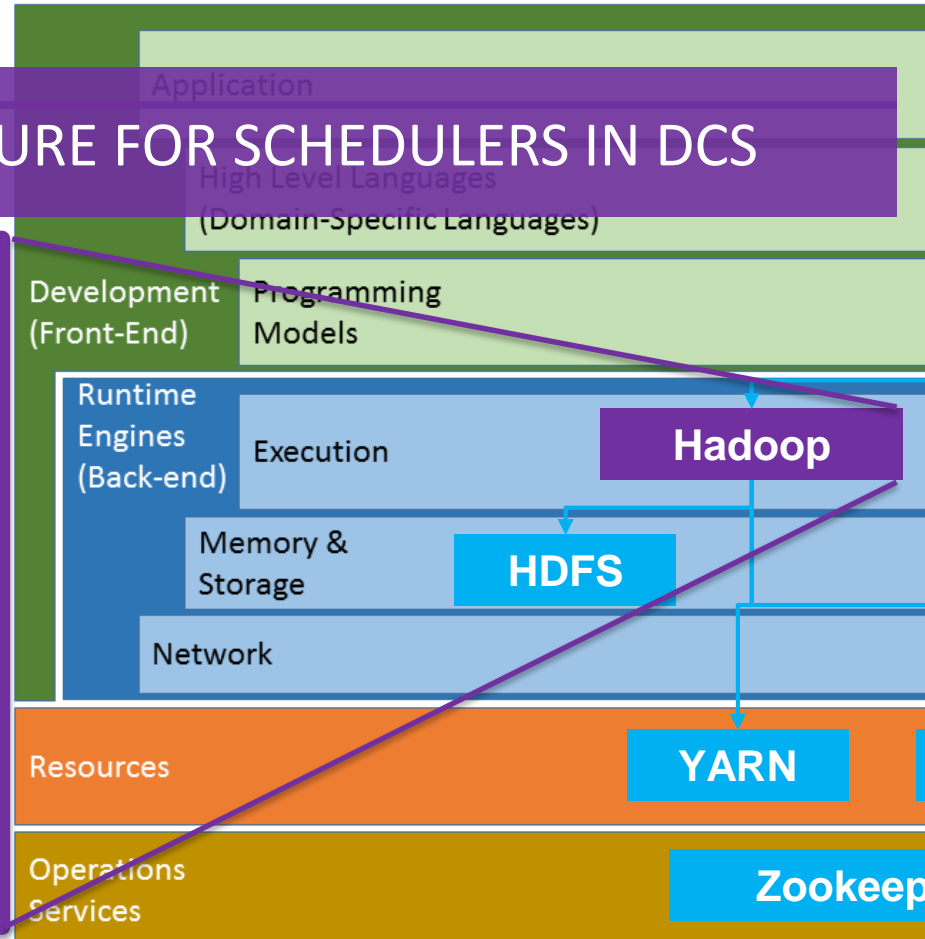
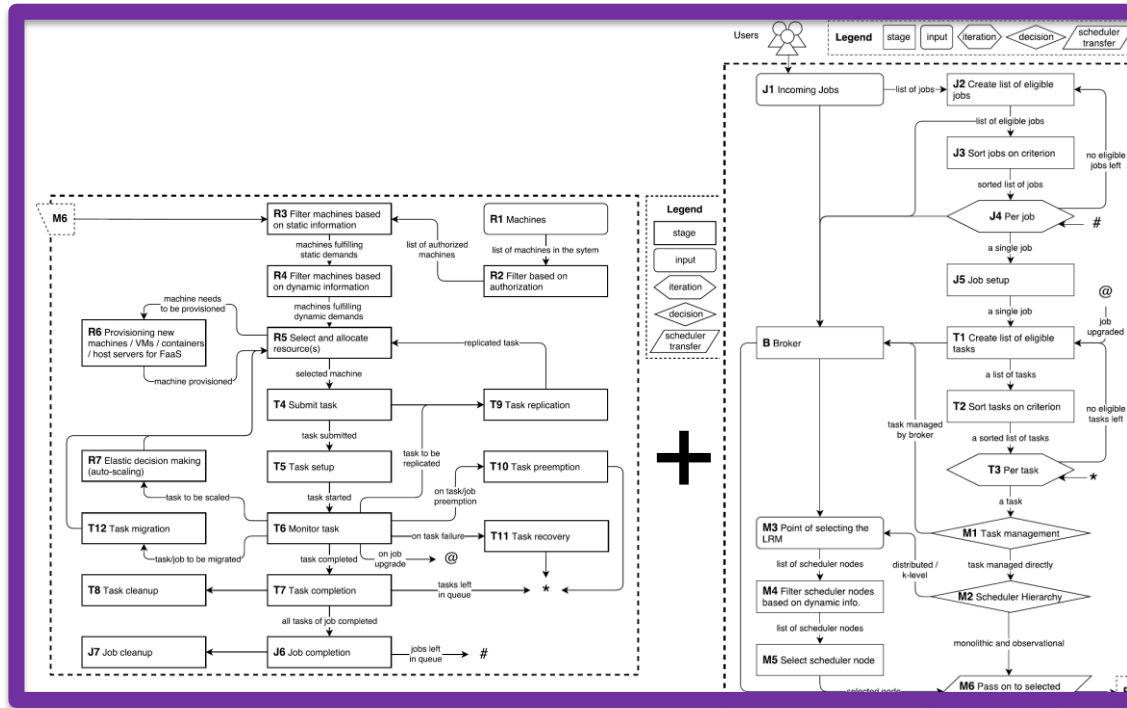
THE COMPLEXITY CHALLENGE

IOSUP ET AL. REFERENCE ARCHITECTURE FOR DCS



Georgios Andreadis

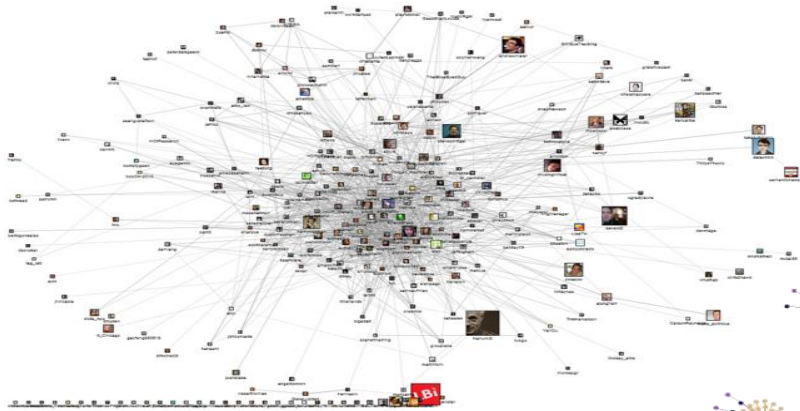
ANDREADIS ET AL. REFERENCE ARCHITECTURE FOR SCHEDULERS IN DCS



[Andreadis et al. SC'18]


MEANINGFUL DISCOVERY

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




Social network

~1 billion vertices
~100 billion
connections



Web graph

~50 billion pages
~1 trillion hyperlinks

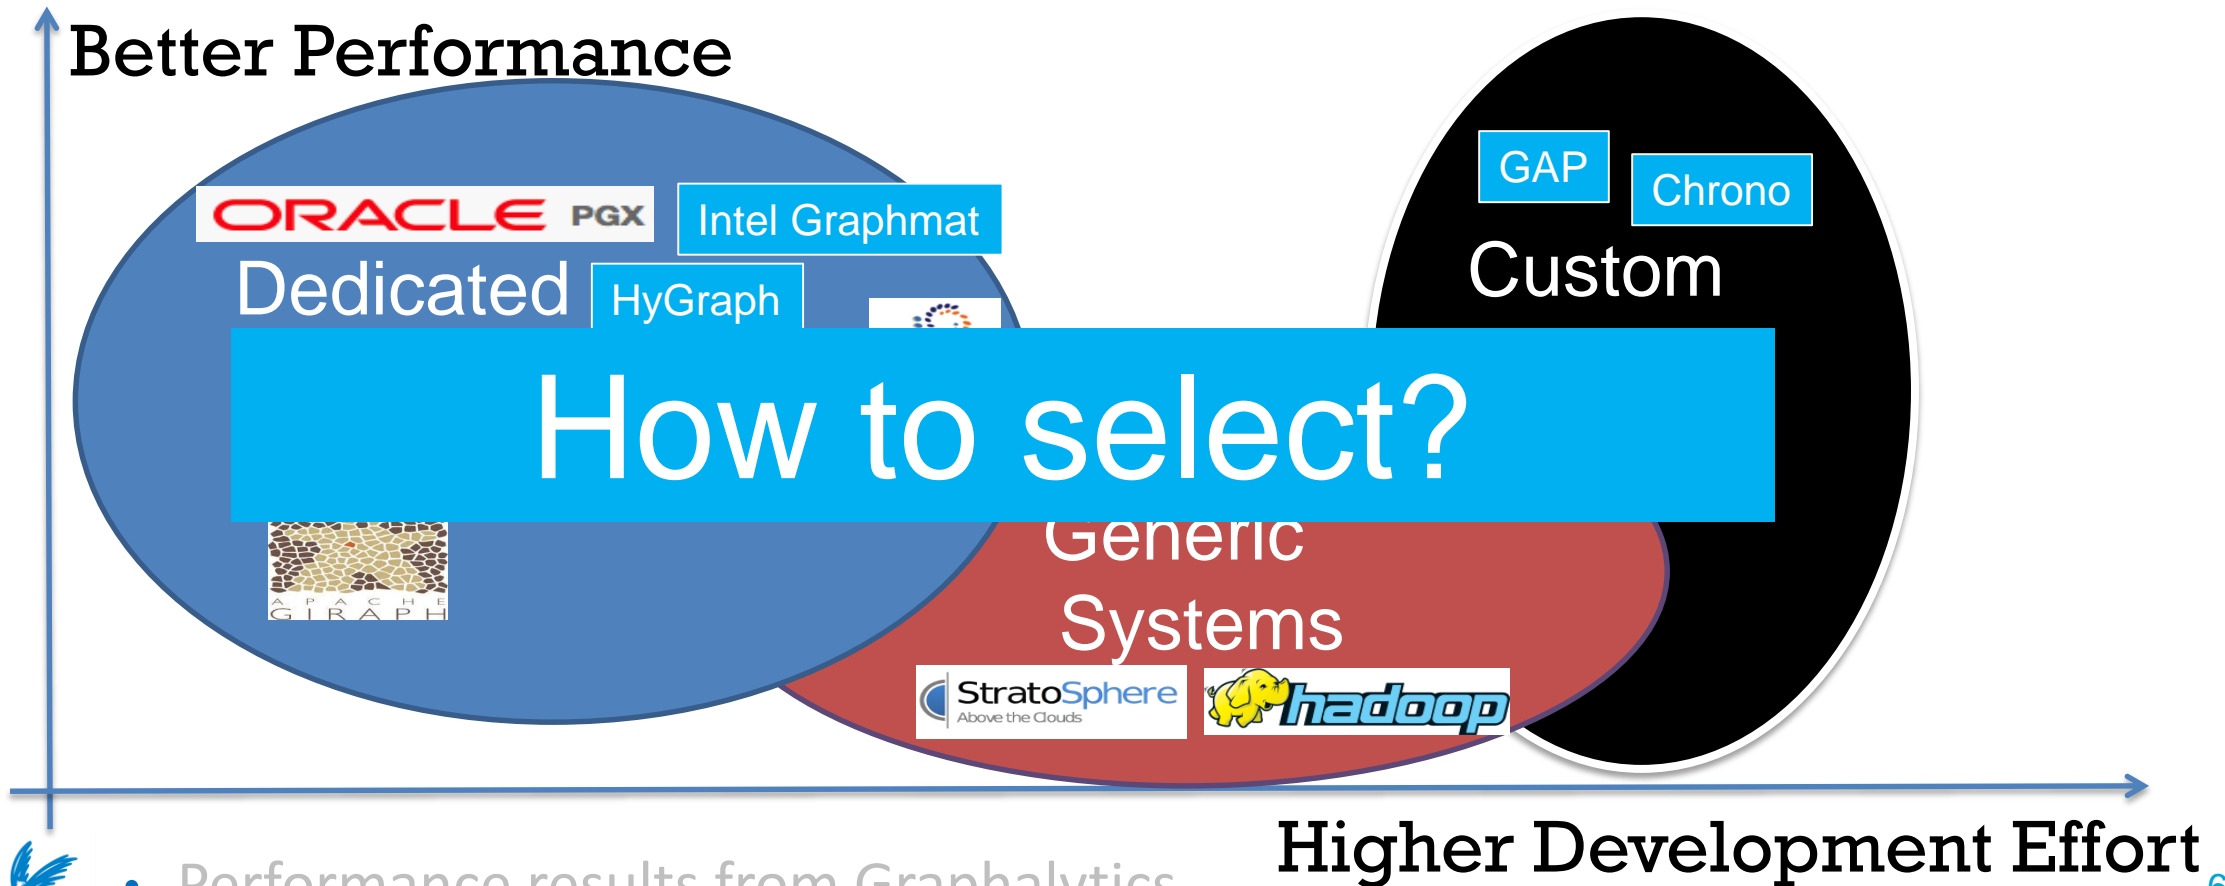


Brain network

~100 billion neurons
~100 trillion
connections

MEANINGFUL DISCOVERY

ENGINEERING LDBC GRAPHALYTICS: THE SYSTEMS LANDSCAPE



MEANINGFUL DISCOVERY

ENGINEERING LDBC GRAPHALYTICS: BENCHMARKING LEADING TO SCIENCE




- **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




MEANINGFUL DISCOVERY

ENGINEERING LDBC GRAPHALYTICS: BENCHMARKING LEADING TO SCIENCE



- **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



- Community endorsed:

graphalytics.org

- Surprising findings:

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

- Triggered new research



MEANINGFUL DISCOVERY

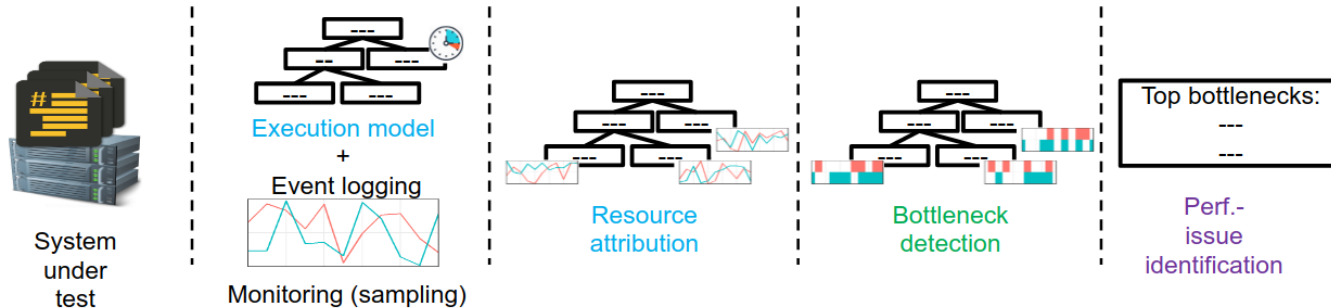
ENGINEERING LDBC GRAPHALYTICS: MODELING LEADS TO PERFORMANCE ANALYSIS



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



Tim
Hegeman



Multi-stage process,
works in ecosystem




MEANINGFUL DISCOVERY

ENGINEERING LDBC GRAPHALYTICS: MODELING LEADS TO PERFORMANCE ANALYSIS



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



- Without Grade10:

No bottleneck at all

- With Grade10:

Always bottleneck

Cause:

- + Message queue full
- + Garbage collector
- + CPU
- + Others



System
under
test




Monitoring (sampling)



Resource
attribution



Bottleneck
detection



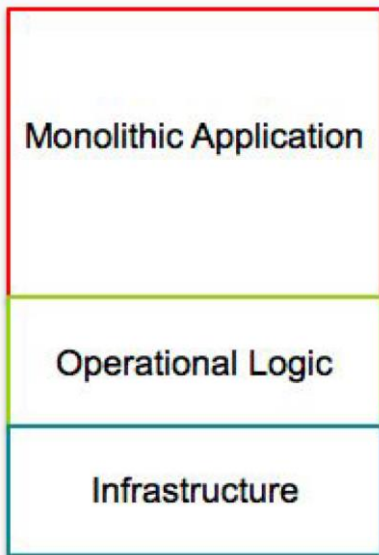
Perf.-
issue
identification

Multi-stage process,
works in ecosystem

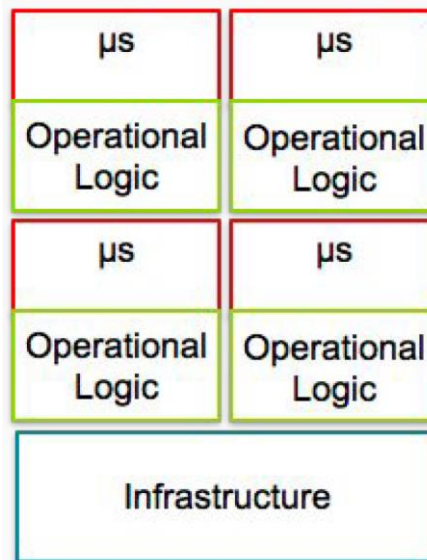
MEANINGFUL DISCOVERY

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

