

Massivizing Computer

Systems = Making Computer

Systems Scalable, Reliable,

High-Performance, etc., Yet

Form an Efficient Ecosystem

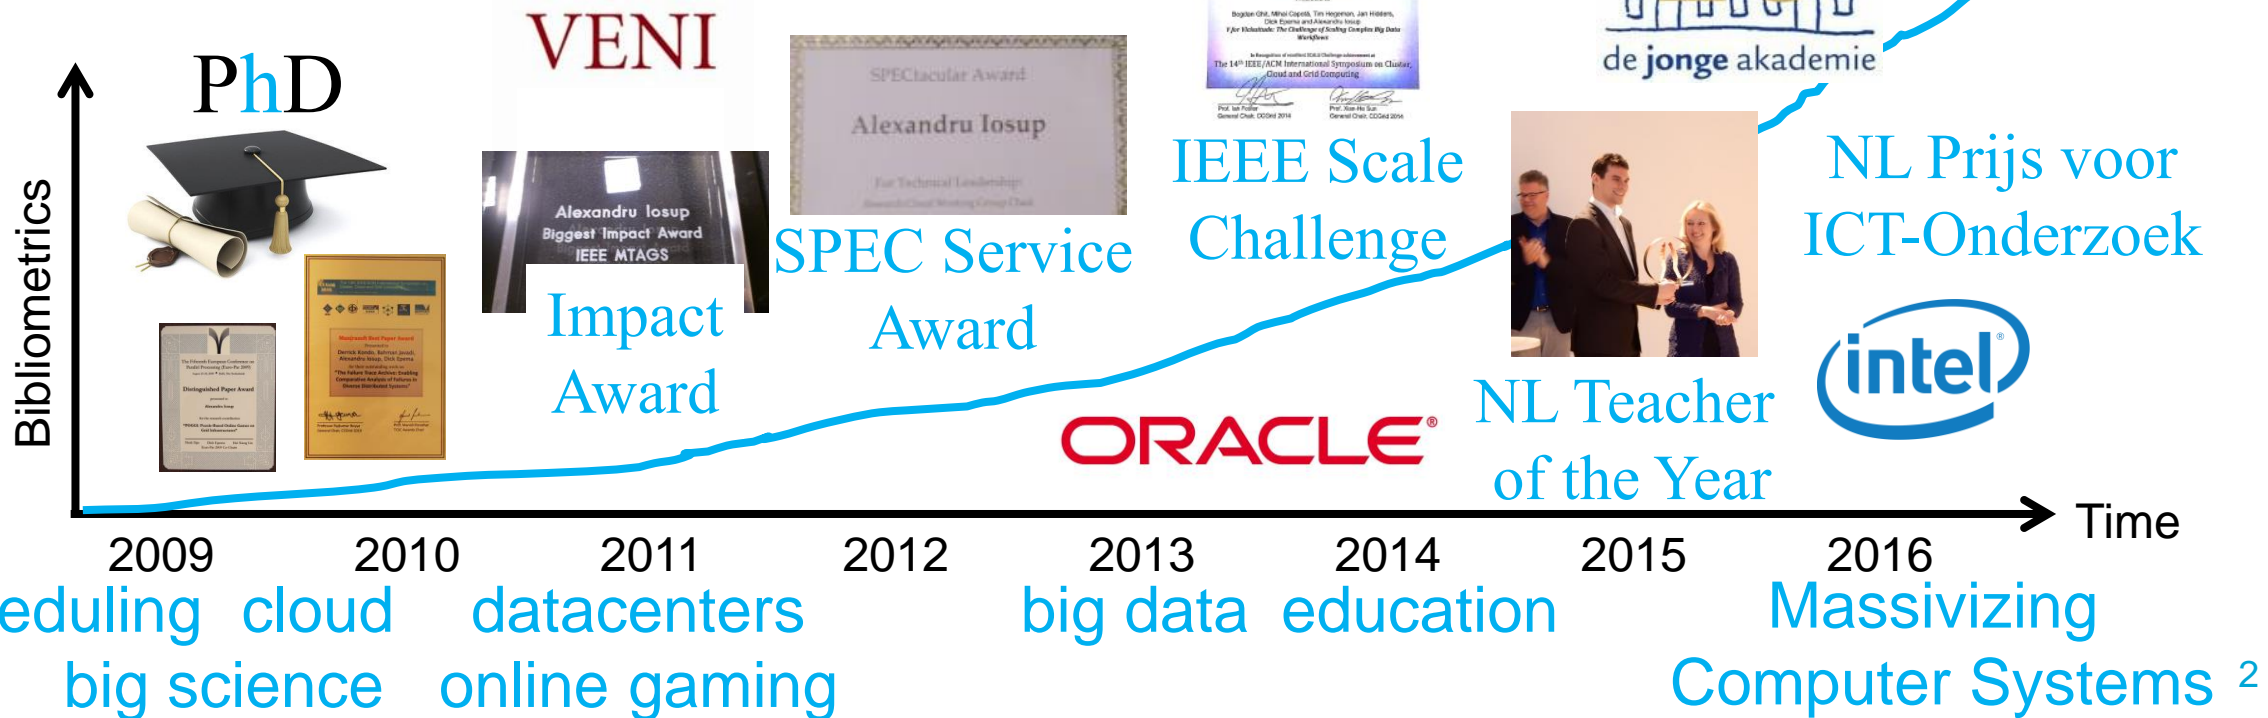


prof. dr. ir. Alexandru



My Background

1. Numerous **award wins** and nominations in **research, education, and service**.
2. **Leadership** in organizing top-level conferences in comp.sci.: HPDC, ICPE, Cluster.
3. Young **Full Professor** (age 36) and **University Research Chair at Vrije Universiteit**, in my field, world top 5-10% under 40, by external reviewers (2016), #1 under 40, by national reviewers (2016), #3 in the world Rising Star (2011).





Dick



Henk





Johan



Fernando

Many Thanks to Many!

<http://atlarge-research.com>



Boxun



Yon



Elvan



Nezh



Vincent



Boddar



Otto



Mihai



Tim



Srai



Jie



Orna



Adele



StefanH



Corina



AlexO



Jesse



Wing



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

This Is the Golden Age of Computer Systems




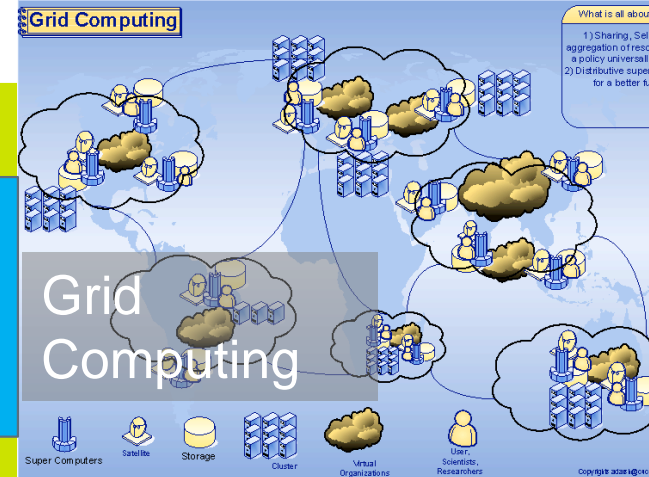
Education for Everyone (Online)



Business Services



cloud Computing




Big Science

Here is how this works...



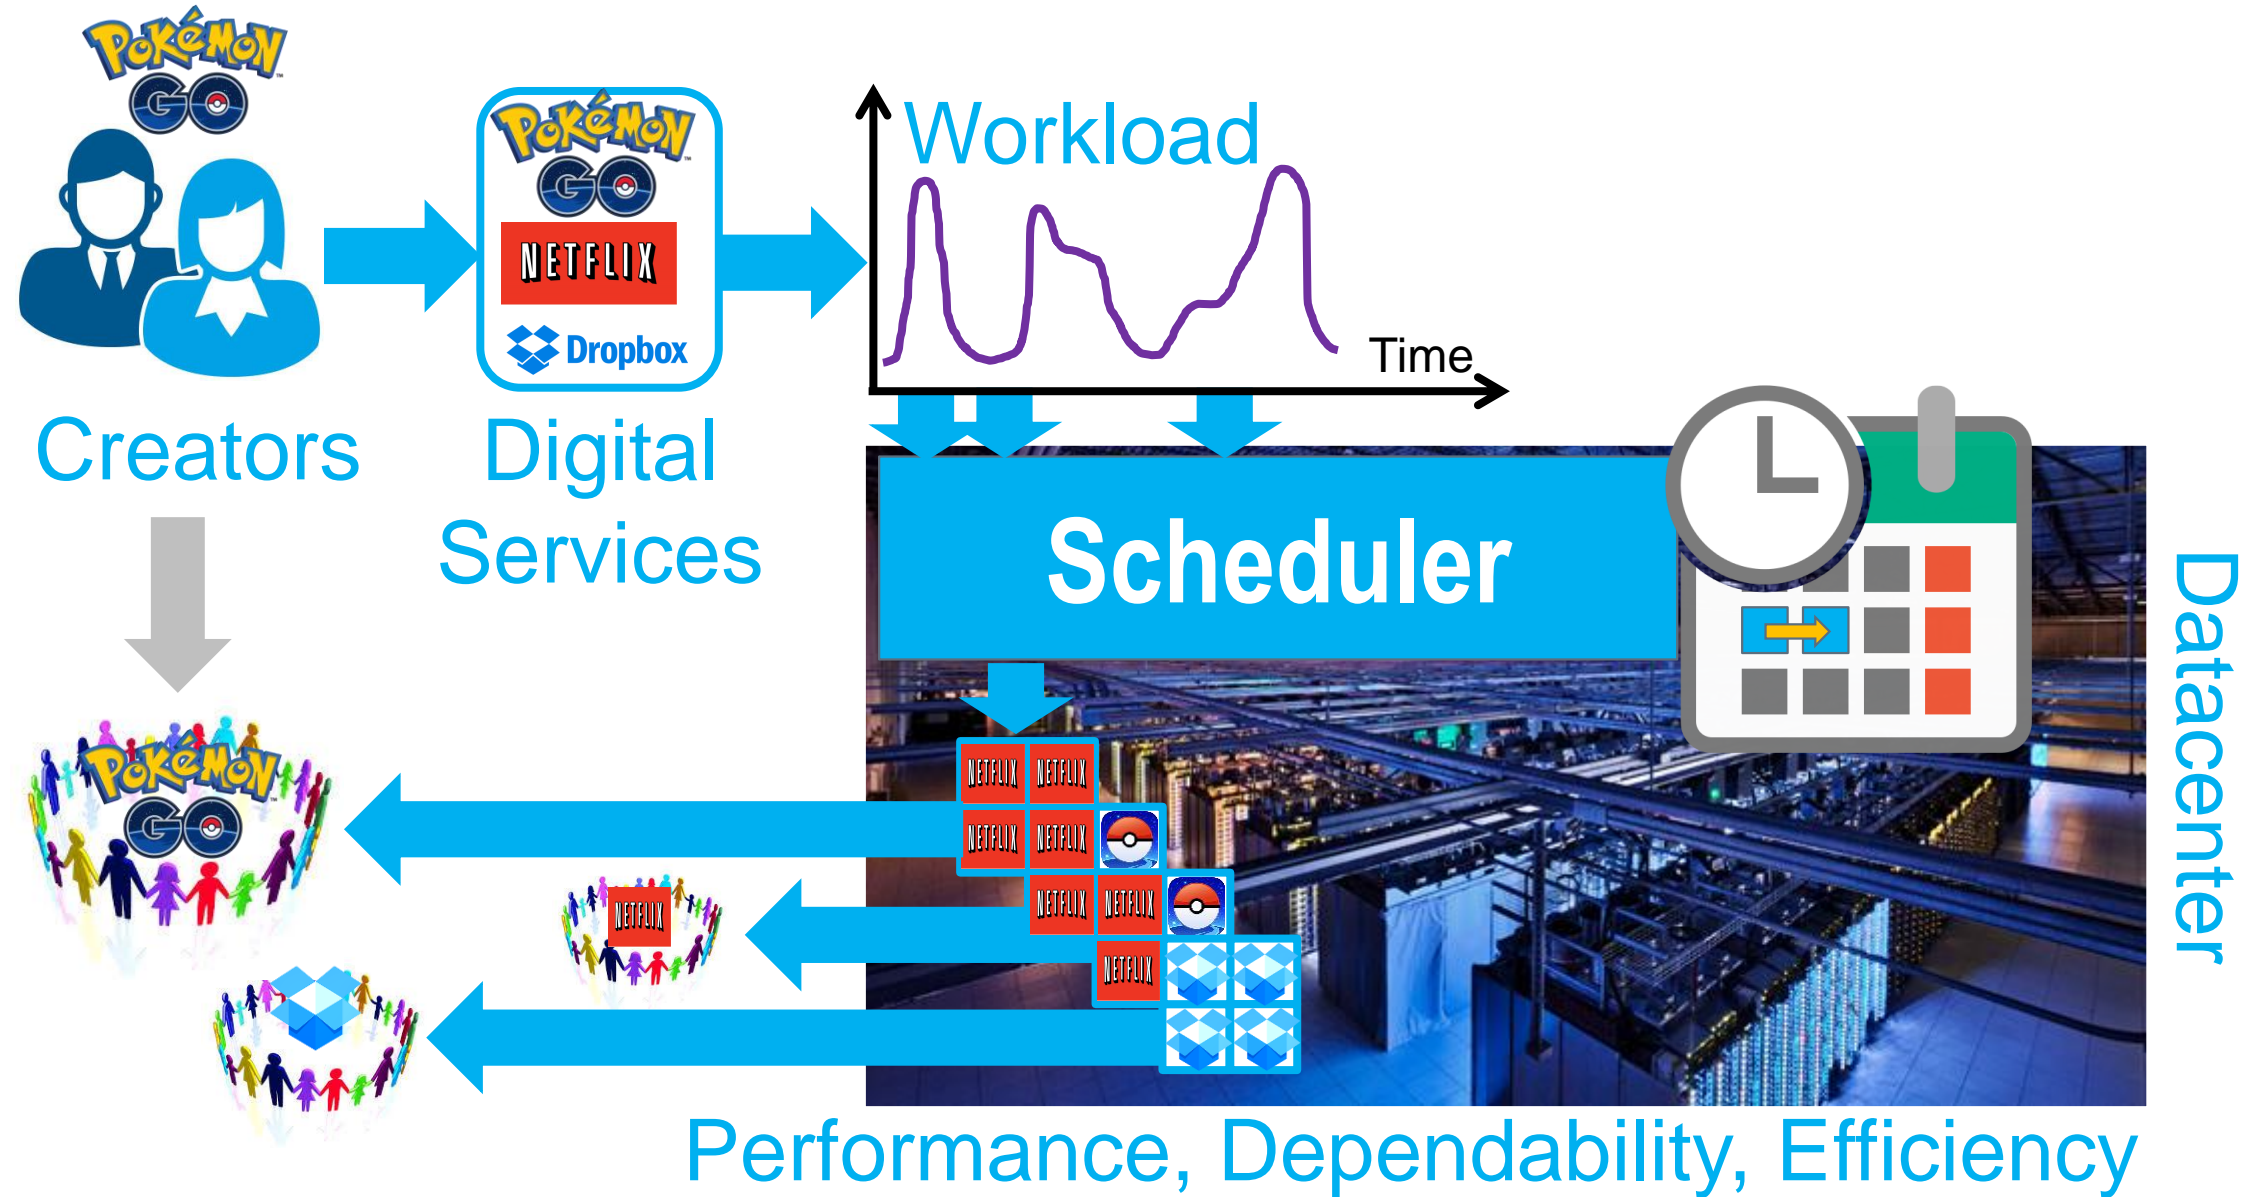
Daily Life



Datacenters



Current Technology: Scheduler? Datacenter? Etc.



(full-/micro-/nano-)

The Golden Age of Computer Systems

... Yet We Are in a Crisis




Education for Everyone (Online)



Business Services



cloud Computing




Big Science

A Crisis? What crisis?!



CTAAAGATGATCTTTAGTCCCGTTGAA
TCTTTAGTCCCGTTGATAACCAACCC
GTAATACCAACCGGACATAAGATCCGG
GGGACTAAAGTCCACCCCTATATATG

TTCAAAATTTCTCAAAAAGAGGGGAG
GTGATTACATACAAATGGAGGTGCCTA
TTTGTACACTACATTTGCACCTATGTTT
GTAAGTTGATGAGAGAAAATGTGTGT



Online Gaming



My other computer is a data center

BIG DATA



Datacenters



ABN·AMRO

Daily Life



AVERAGE DAILY ONLINE GAMERS WORLDWIDE

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



The Crisis: In the Digital Economy, Few Can Afford Being Successful!

Why does this happen?

What to do about it?

My Research

NETFLIX

Dropbox

CERN

“ICT is vital for SMEs, SMEs are 60% GDP”
“15% ICT market is simple cloud services”
“Already 60+ bn.€/year”

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

The Scheduling Challenge

**“30—70% scheduler decisions
incorrect in datacenters”**

Source: IEEE Computer'15

**“current schedulers not efficient
for many users, diverse services”**

Source: Dutch industry, CCGRID'15

**“new schedulers not used in
datacenters, fear of failure”**

Source: EuroPar'13,'14



Need Smarter Schedulers



Need to Select Schedulers

The Dependability* Challenge

* Availability, Reliability, etc.

The Register
Biting the hand that feeds IT

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

Need Dependable Systems

THE VERGE

TRENDING NOW

The new Nvidia Shield is the 'world's first 4K Android TV console' and launches this May for \$199...

26
NEW ARTICLES

LOG IN | SIGN UP | LONGFORM | VIDEO | REVIEWS | TECH | SCIENCE | ENTERTAINMENT | DESIGN | BUSINESS | US & WORLD | FORUMS

APPS | TECH

www.theverge.com/2014/2/23/5439398/whatsapp-founder-apologizes-for-our-longest-and-biggest-outage-in

WhatsApp founder apologizes for 'our longest and biggest outage in years'

82
COMMENTS




By [Russell Brandom](#) on February 23, 2014 12:25 pm | [Email](#) | [@russellbrandom](#)

DON'T MISS STORIES [FOLLOW THE VERGE](#) | [Like](#) | [Follow](#) | [Subscribe](#) | [Follow](#)

The New World Challenge



Cloud operator: new value-adding services, new workloads including FaaS, DevOps workloads



Need Operational Models

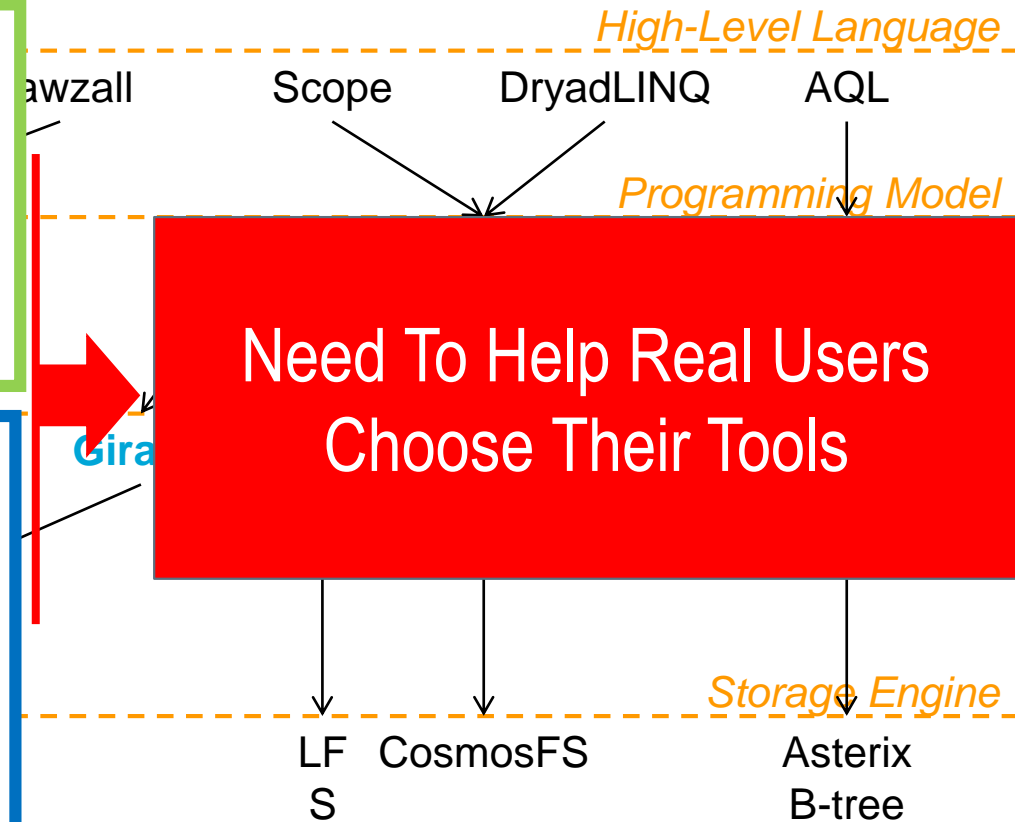
Cloud customer: new apps, new services, micro-services, customers can become operators (value-chain)



The Ecosystem Navigation Challenge

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

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



Jevons Effect: More Efficient, Yet Less Capable

Nov 2015: Over 500 YouTube videos have at least 100,000,000 viewers each.

Jun 2017: How many are there?

If you want to help kill the planet:

https://www.youtube.com/playlist?list=PLirAqAtl_h2r5g8xGajEwdXd3x1s

**Need To Be Much More
Efficient, But Also To
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 billion views (Nov 2015).

The New “Jevons Effect”: The “Data Deluge” Challenge



Data Deluge =

Need To Address The
“Data Deluge”

To be capable of processing Big Data, need to address Volume, Velocity, Variety of Big Data*

* Other Vs possible: ours is “vicissitude”

This Is the Golden Age of Computer Systems and We Have Many Tools... Yet We Are in a Crisis

Need to Understand How to Use Our Tools

... but the Current Laws and Theories Were Built For Isolated Computer Systems

Need Smarter Schedulers

Need to Understand Operational Laws when Massivizing Computer Systems

Need Dependable Systems

Need to Create Theories on how to Massivize Computer Systems while Ensuring Wanted Properties

Need to Address “Data Deluge”, “Ecosystem Navi”, etc.

Need to Be Much More Efficient, But Also Ethical

Need to Build, to Massivize Computer Systems with Wanted Properties

This Is the Golden Age of Computer Systems ... Yet We Are in a Crisis

Massivizing Computer Systems
Tackles All These Challenges...


... and Is Relevant, Impactful, and
Inspiring for Many Young Scientists

Massivizing Computer Systems

In Pasteur's Quadrant+:

- Fundamental research
- Inspired by real use
- Experimental in nature

~ Big Science as management,
including int'l. collaborations



Fundamental Research in Massivizing Comp. Sys.

Scheduling

Bags-Of-Tasks

Workflows

Portfolio

Dependability

Failure Analysis*

Space-/Time-Correlation

Availability-On-Demand

New World+

Workload Modeling

Business-Critical

Online Gaming

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

+ Please ask for a definition

My Contribution So Far Personal grants

* Award-level



An Example: Portfolio Scheduling for Datacenters (what's in a name)



Massivizing Computer Systems

Scheduling

Bags-Of-Tasks

Workflow

Mixed-Workload

Portfolio

Dependability

Failure Analysis*

Space-/Time-Correlation

Availability-On-Demand

New World

Workload Modeling

Interaction Graphs

Business-Critical

Online Gaming

Possible topic for collaboration?

Ecosystem Navigation

Performance Variability

Grid*, Cloud, Big Data

Benchmarking

Longitudinal Studies

Socially Aware Techniques

Collaborative Downloads*

Groups in Online Gaming

Toxicity Detection*

Software Artifacts

Graphalytics, etc.



Data Artifacts


Distributed Systems Memex*

Fundamental Problems

My Contribution So Far (* Award-winning)

Portfolio Scheduling, In A Nutshell

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

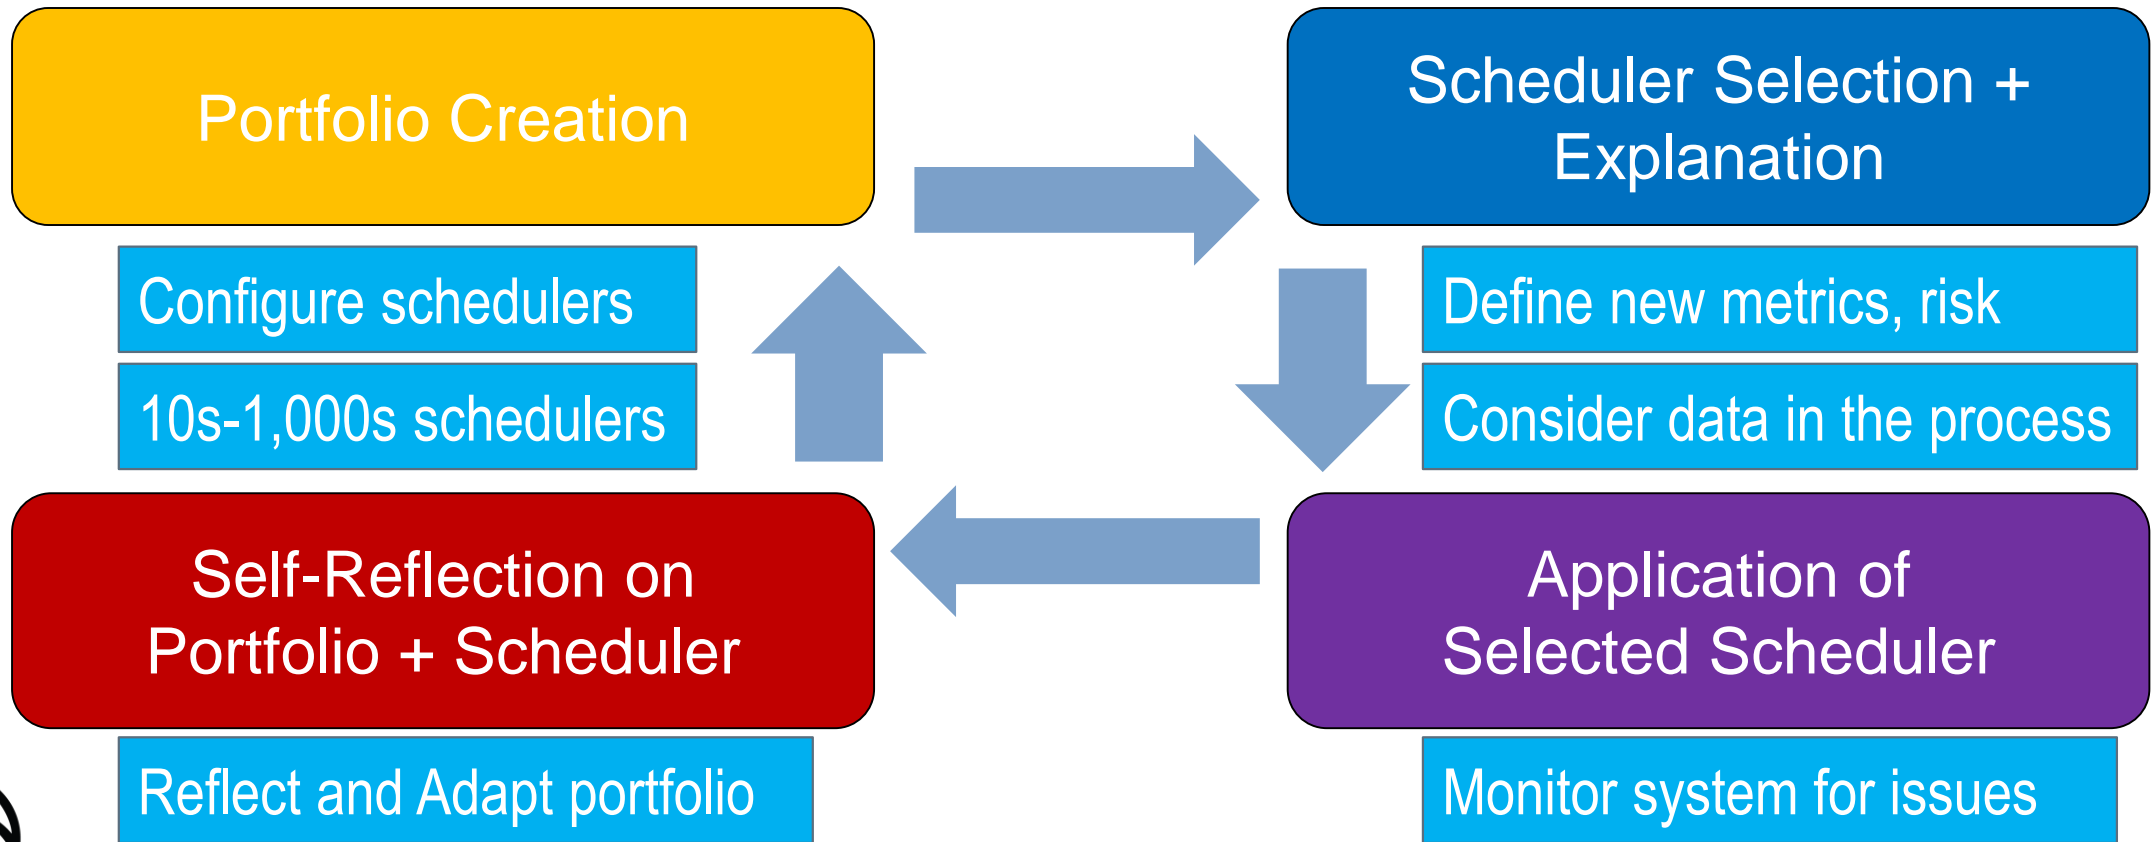


1. Create a set of schedulers
 - Resource provisioning and allocation policies for datacenters
2. Select active scheduler online, apply for the next period, analyze results
(Repeat)



Portfolio Scheduling for Computer Systems

Portfolio Scheduling





Experimental Research Methodology

My Main Scientific Instrument: DAS-5

My (& Your) Prototypes



SURFnet6

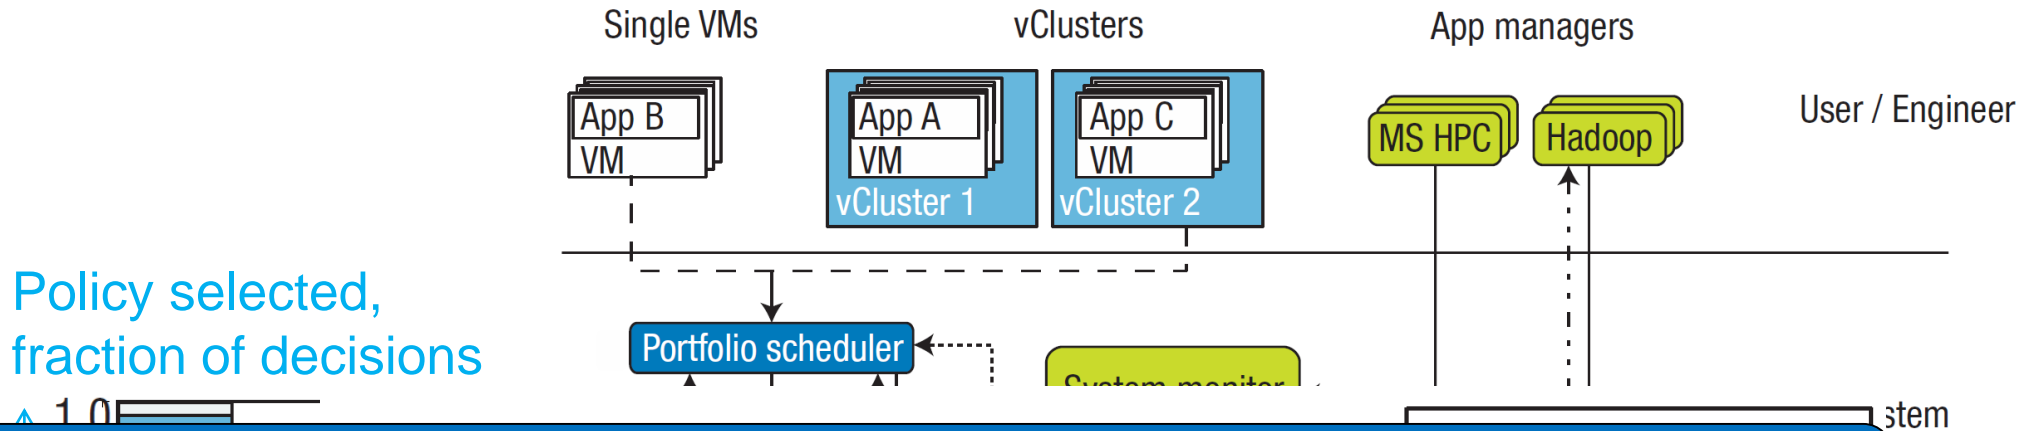


300+ scientists as users

We won IEEE Scale Challenge 2014




Portfolio Scheduling in Practice: Massive 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)



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

- ↔ Scenario
- DC w System components
- ← System/Datacenter components
- ← VM profile data
- ⋯ API Commands
- ← Monitoring data
- ↔ Nebu application connection



End of Example:
Portfolio Scheduling for Datacenters
(what's in a name)





Take-Home Message: We Are Massivizing Computer Systems



- The Golden Age of Computer Systems →
- My Research is about Massivizing Computer Systems →
- **General Questions** ← we are here now



1. Students involved in this topic are successful.
2. Computer Systems: both the Golden Age and in crisis.
3. Big Science as management structure.
4. Interested in international collaboration.
5. Ask for details: Theory? → Problem in Pasteur's quadrant? →



Education for Everyone



Business Services



Massivizing Computer Systems



Big Science

Consider Reading the Following:

1. Iosup et al. LDBC Graphalytics: A Benchmark for Large-Scale Graph Analysis on Parallel and Distributed Platforms. PVLDB 9(13): 1317-1328 (2016)
2. Guo et al.: Design and Experimental Evaluation of Distributed Heterogeneous Graph-Processing Systems. CCGrid 2016: 203-212
3. van Beek et al.: Self-Expressive Management of Business-Critical Workloads in Virtualized Datacenters. IEEE Computer 48(7): 46-54 (2015)
4. Jia et al.: Socializing by Gaming: Revealing Social Relationships in Multiplayer Online Games. TKDD 10(2): 11 (2015)
5. Ghit et al.: V for Vicissitude: The Challenge of Scaling Complex Big Data Workflows. CCGRID 2014: 927-932
6. Guo et al.: How Well Do Graph-Processing Platforms Perform? An Empirical Performance Evaluation and Analysis. IPDPS 2014: 395-404
7. Javadi et al.: The Failure Trace Archive: Enabling the comparison of failure measurements and models of distributed systems. J. Parallel Distrib. Comput. 73(8): 1208-1223 (2013)
8. Iosup and Epema: Grid Computing Workloads. IEEE Internet Computing 15(2): 19-26 (2011)
9. Iosup et al.: On the Performance Variability of Production Cloud Services. CCGRID 2011: 104-113
10. Iosup et al.: Performance Analysis of Cloud Computing Services for Many-Tasks Scientific Computing. IEEE Trans. Parallel Distrib. Syst. 22(6): 931-945 (2011)