


An Introduction to #CloudComputing



Tim
Hegeman



Georgios
Andreadis

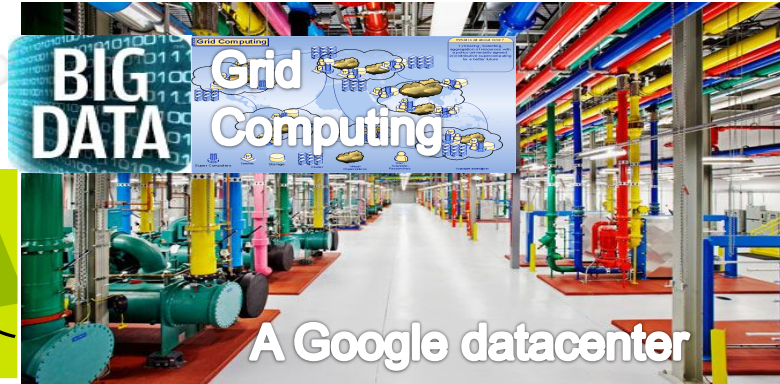



Fabian
Mastenbroek




Prof. dr. ir. Alexandru Iosup

ICT Addresses Major Societal Challenges



The quadruple helix: prosperous society & blooming economy & inventive academia & wise governance depend on ICT

- Enable data access & processing as a fundamental right in Europe
- Enable big science and engineering (2020: €100 bn., 1 mil. jobs)
- “To out-compute is to out-compete”, but with energy footprint <5%
- Keep Internet-services affordable yet high quality in Europe
- The Schiphol of computation: Netherlands as a world-wide ICT hub



Societal Challenges, A Concrete Example: Graph Processing for Everyone

LinkedIn

Oracle 1.2M followers,
132k employees

company/day:
40-60 posts, 500-700 comments

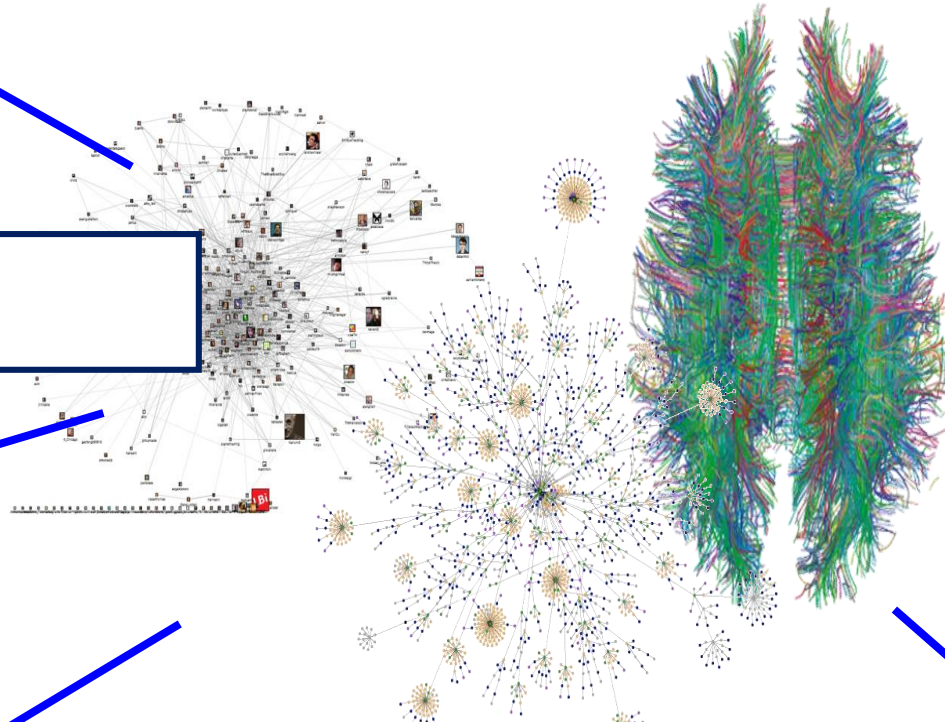


270M MAU
200+ avg followers
>54B edges



1.2B MAU 0.8B DAU
200+ avg followers
>240B edges

YAHOO!



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

friendster
(SME*)

XFIRE
(SME*)

* SMEs in EU/NL= 60% gross value added, little to no ICT expertise

LinkedIn: Job-Seeker & Job-Creator Match

The State of LinkedIn



Q: How can LinkedIn offer this service to you?

**Introducing cloud computing
("cloud", not "Cloud")**

reminder of not only where we've been, but also where we're headed as we work to create economic opportunity for every professional in the world.

What is Cloud Computing?

1. A Cloudy Buzzword

- 18 definitions in computer science (ECIS'10). NIST has one. Cal has one. We have one.
- “We have redefined cloud computing to include **everything that we already do.**” Larry Ellison, Oracle, 2009



What is Cloud Computing?

2. A Descendant* of the Grid Idea

* Subset.



Source: <http://royal.pingdom.com/2008/04/11/map-of-all-google-data-center-locations/>

“A computational grid is a hardware and software infrastructure that provides dependable, consistent, pervasive, and inexpensive access to high-end computational capabilities [+ for] nontrivial QoS.” I. Foster, 1998 + 1999

Cloud MW Stack

~~Cloud Grid~~ Applications

~~Cloud Grid~~ Very High Level MW

~~Cloud Grid~~ High Level MW

~~Cloud Grid~~ Low Level MW

Virtualized HW + OS

MW = Middleware

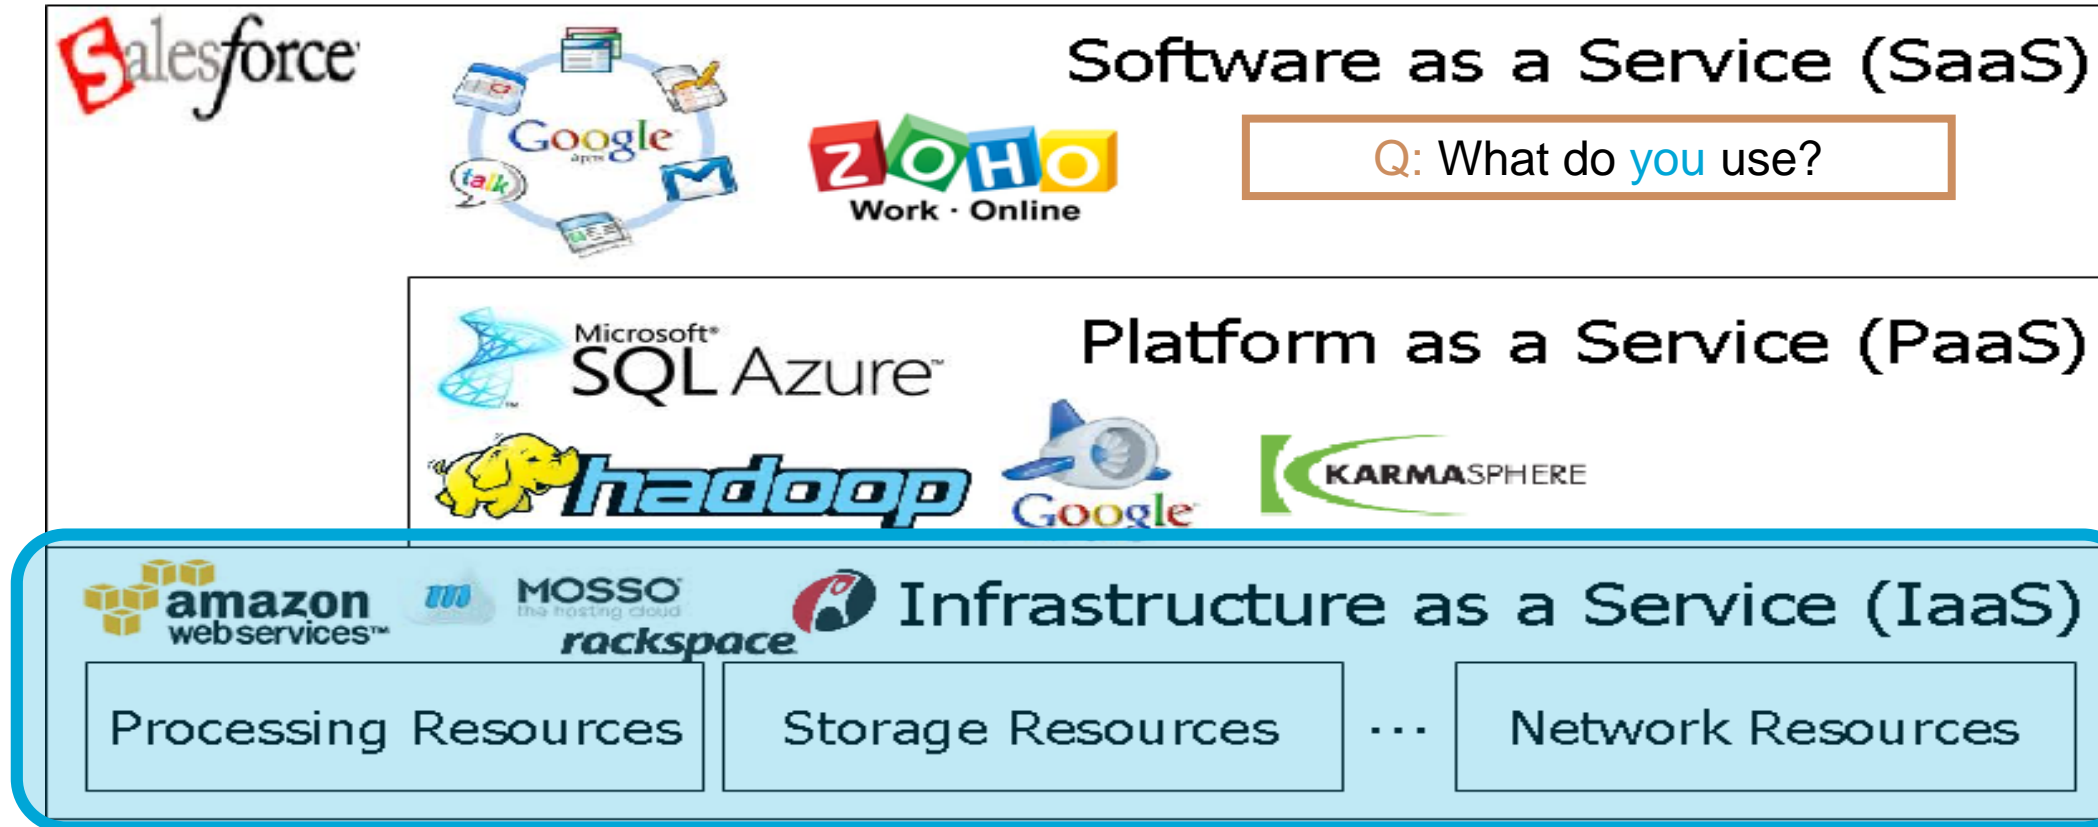


Have you noticed “QoS”? What is that?

What is Cloud Computing?

3. A Useful IT Service

“Use only when you want! Pay only for what you use!”



Main Characteristics of IaaS Clouds

1. On-Demand Pay-per-Use
2. Elasticity (cloud concept of Scalability)
3. Resource Pooling
4. Fully automated IT services
5. Quality of Service

Q: Sounds great, but ...
How can we make
all this stuff happen?

**Introducing datacenters &
datacenter-based clouds =
ICT service creation
for everyone**



Intro to Cloud Computing

TU: 1 — Pitch on Datacenter-Based Cloud Computing

1 — The Golden Age of Datacenters

4 — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

7 — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

— Reality Check

Here or @home

Interactive

This Is the Golden Age of Datacenters



Personalized & Online Education





Business Services




cloud Computing

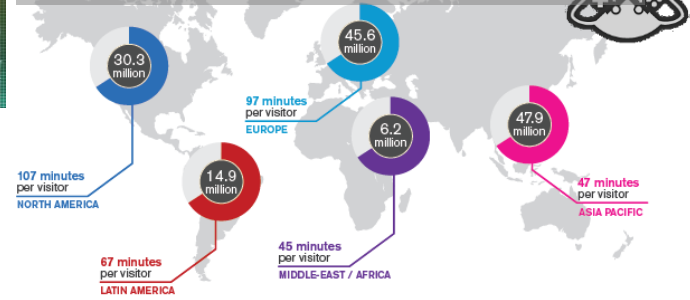
Grid Computing



Big Science



Online Gaming



My other computer is a data center



AVERAGE DAILY ONLINE GAMERS WORLDWIDE

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



All rights reserved



100% Green

Factories Powering the Goods Economy and Better Living Standards



Creator



Factory (1980)



Consumer



Datacenters = Digital Factories Powering the Digital Economy and Better Living Standards



Creator



Datacenter (2010)



Consumer

We want more of this

“15% ICT market is simple DC services”

“Already 60+ bn.€/year”

“Consume more energy than the transport industry”

We want less of this (and other issues raised by DCs)



Intro to Cloud Computing



5' — Pitch on Datacenter-Based Cloud Computing

5' — The Golden Age of Datacenters

20' — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35' — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10' — Reality Check

Here or @home

Interactive

Intro to Cloud Computing



5' — Pitch on Datacenter-Based Cloud Computing

5' — The Golden Age of Datacenters

20' — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35' — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10' — Reality Check

Here or @home


Interactive

Joe Has an Idea (\$\$\$)




Solution #1 Buy then Maintain

- Big up-front commitment
- Load variability: **NOT supported**





Solution #2 Deploy on IaaS Cloud



- NO big up-front commitment
- Load variability: supported

Inside a Cloud Datacenter: Infrastructure as a Service

Q: So are we just **shifting the ownership problem**, that is, to the cloud owner?



Intro to Cloud Computing



5' — Pitch on Datacenter-Based Cloud Computing

5' — The Golden Age of Datacenters

20' — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35' — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10' — Reality Check

Here or @home

Interactive

The Cloud Owner Perspective



- Build the datacenter
- Operate the datacenter



Build the datacenter = Servers + Server Racks +
Intra-Rack Network + Inter-Rack Network

An Entire Floor in a
Google Datacenter

Datacenter = commodity high-performance ICT



A Google datacenter

- Large-scale infrastructure
- High-tech automated software to manage
- Inter-connected computer clusters
- High-end computation, storage, network
- Large memory capacity
- “my other computer is a datacenter”

The Pizza-Box Stack

- The 1U server



The Pizza-Box Stack

- The 1U server



The Pizza-Box Stack

- The 1U server



The Pizza-Box Stack

- The 1U server
- The 19" server rack (42U is r
 - Half-racks also common

Q: What is a **half-rack**, and why is it useful?



The Data Center Network



- Network bandwidth per rack
 - 1 x 48 port GigE switch – 40 UP – 8 DOWN links

Q: What is are the characteristics of **research**, **enterprise**, and **consumer** hardware (and other services)?

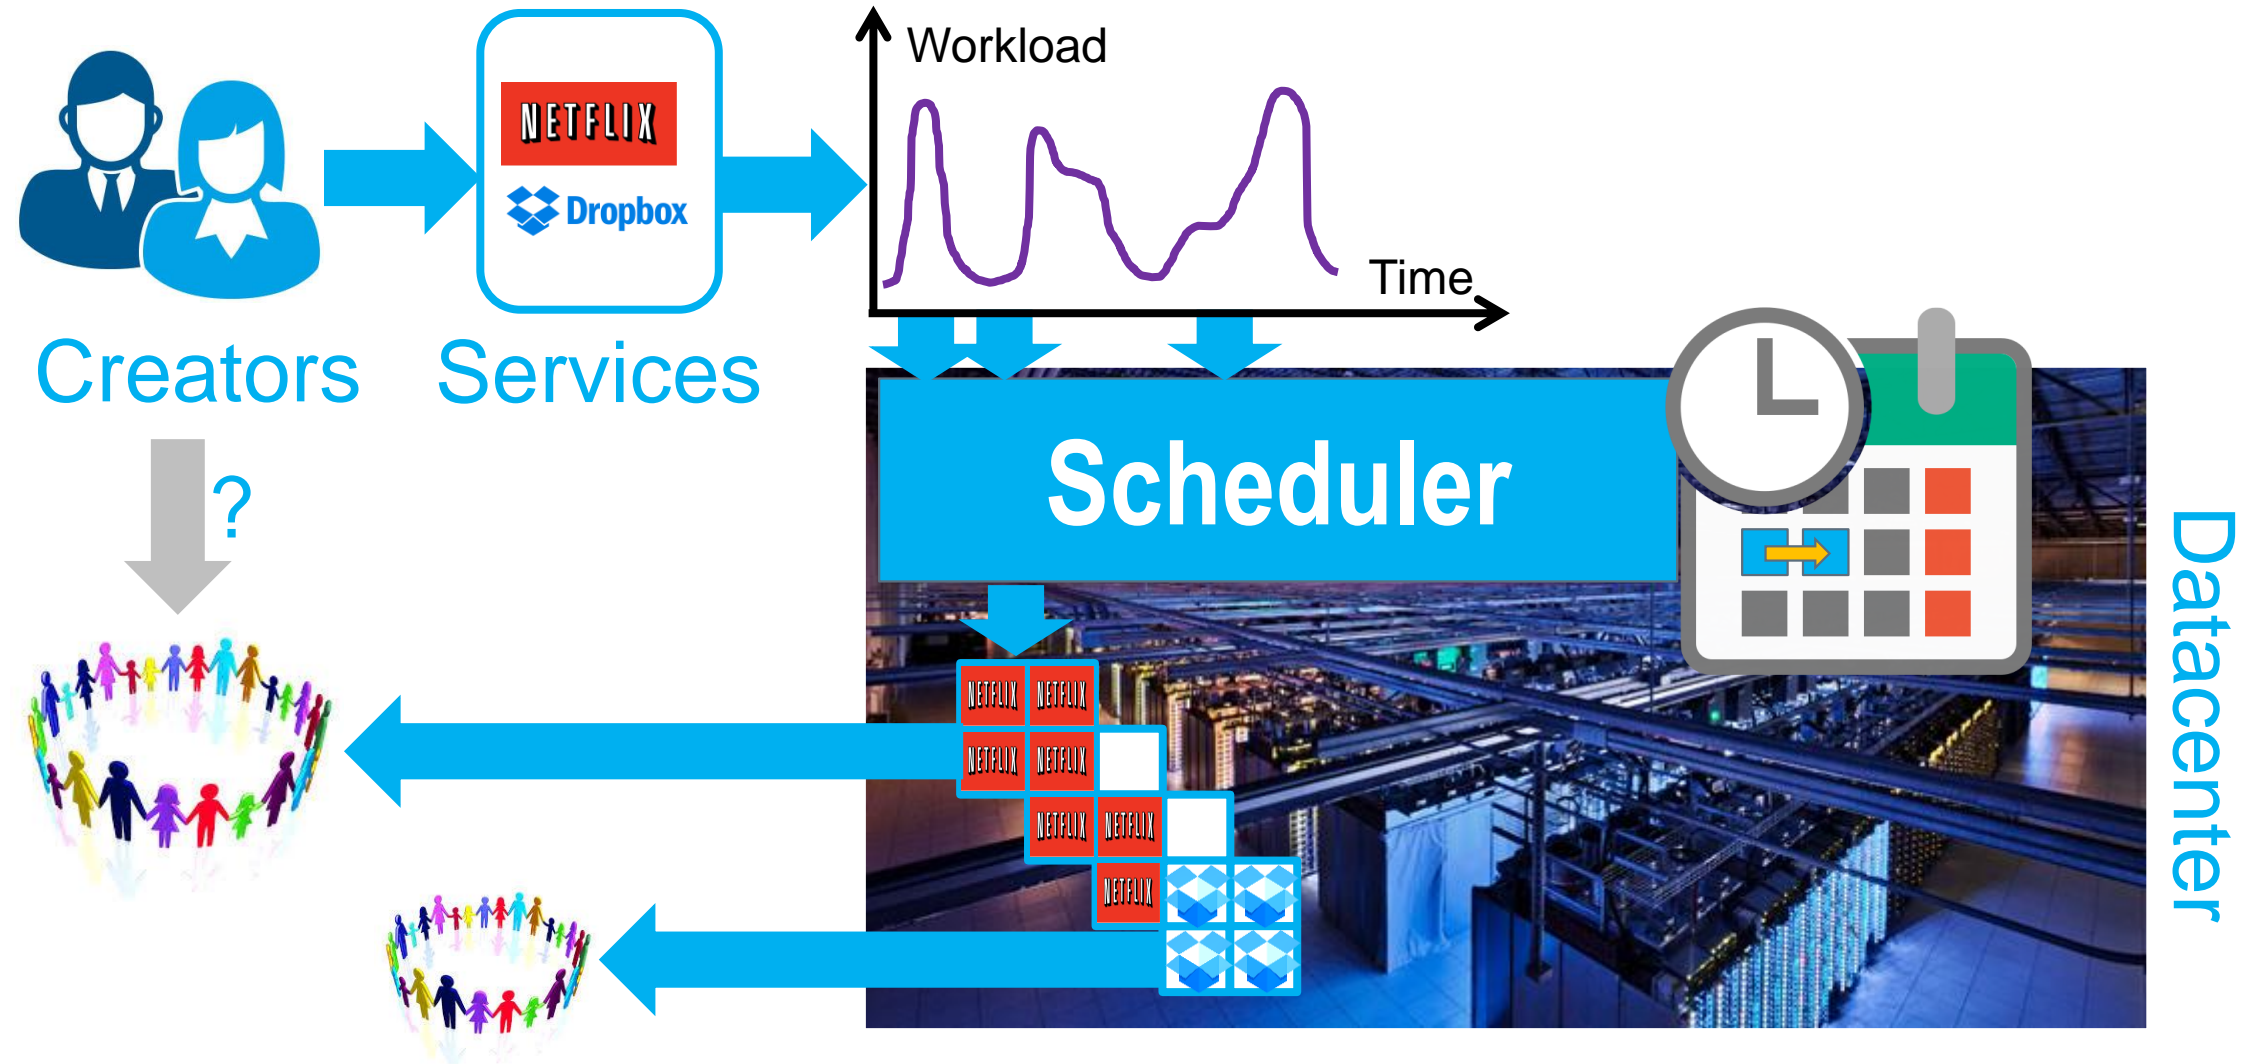
- (consumer) 100 Mbps for 1 GigE rack switch
- (enterprise) 1 Gbps for 10 GigE rack switch
- (enterprise) 10 GBps for ncHT3 (supercomputing class)
- (research) 1 Tbps!~100 GBps optical (not yet production-ready)



What does it mean to operate the DC? Running Workloads




Using a Scheduler to Run Workload



Resource Sharing Models


Grids
Space-Sharing

MusicWave


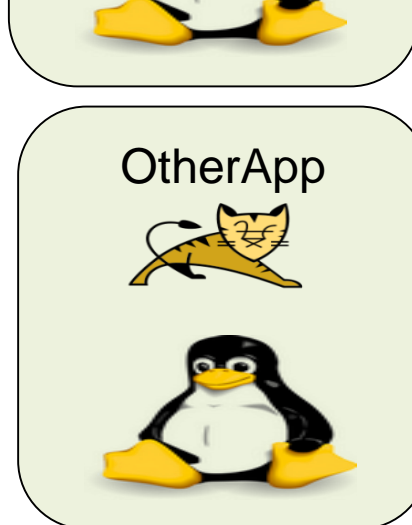


IaaS Clouds
Time-Sharing

Q: Which one is better?



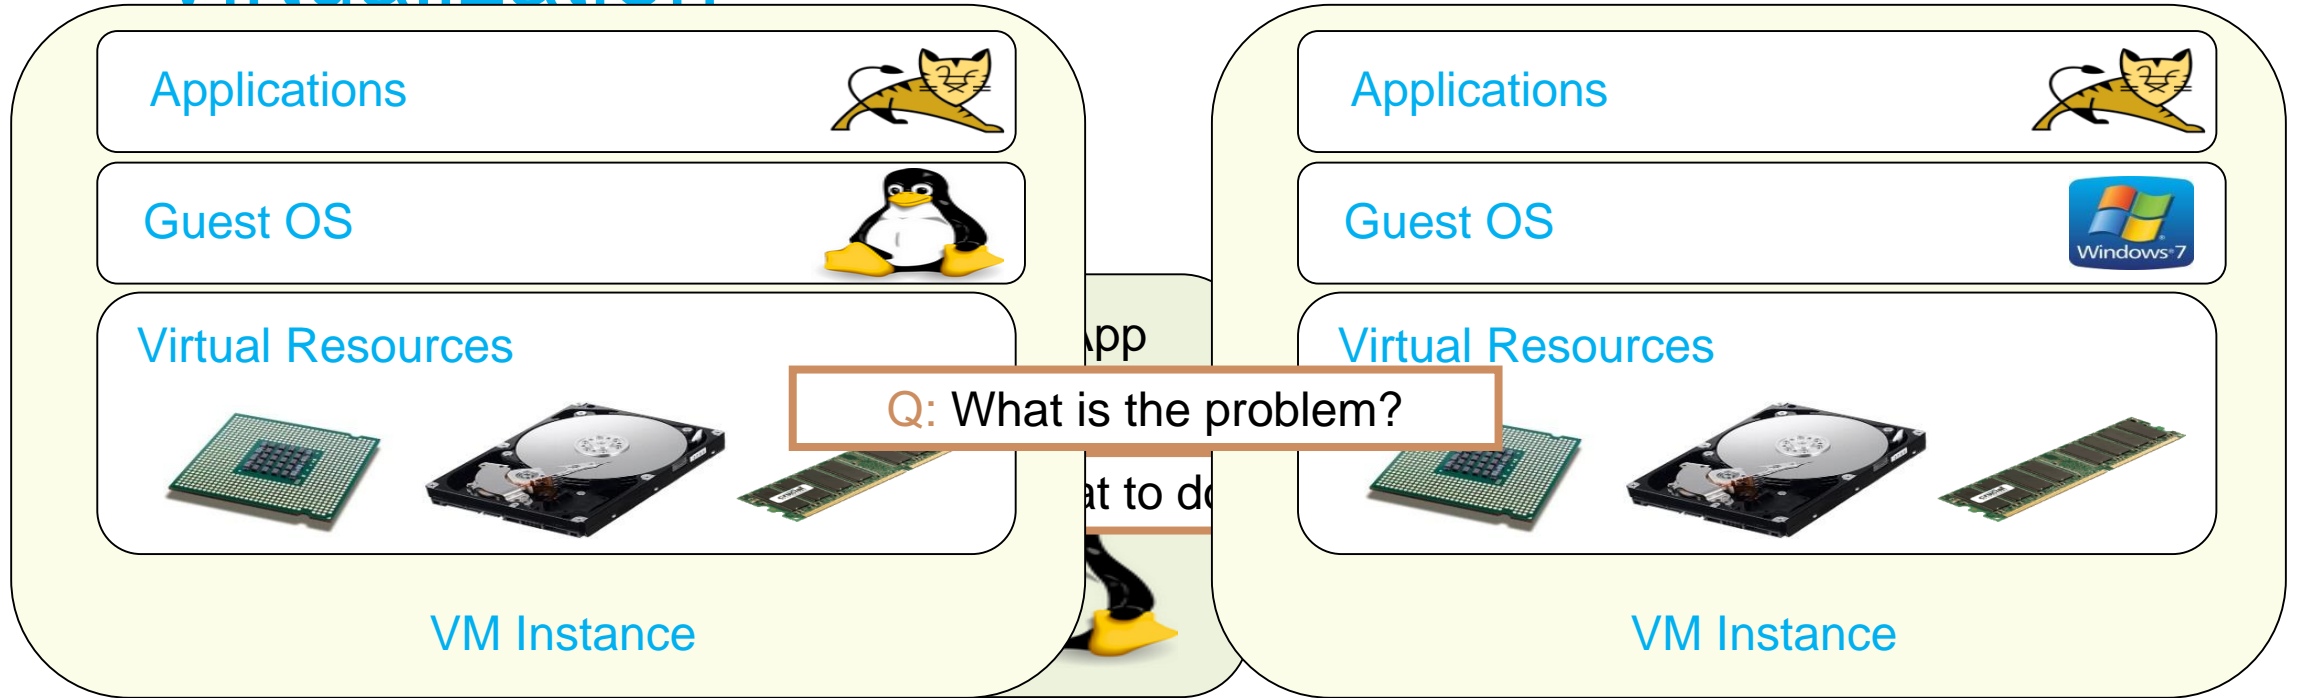
Host OS



Host OS



Virtualization



Virtualization

Host OS

The Cloud User Perspective

- Lease the resources
- Use the resources


Where to Lease Resources From?

IaaS Cloud Deployment Models

Private
On-premises

Public
Off-premises


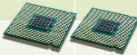
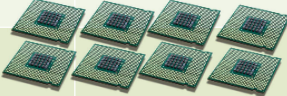
Hybrid
Two or more clouds



Use Case: Amazon Elastic Compute Cloud (EC2)


- Prominent IaaS provider (also Microsoft, Google, Alibaba)
- Datacenters all over the world
- Many VM instance types
- Per-hour **charging**
- **Auto-scaling with simple policies**



Instance	Capacity	US\$/hour
m1.small		0.10
m1.large		0.38
c1.xlarge		0.76

What Does It Mean To Lease A Resource?

The Virtual Machine Lifecycle





At this scale, computer science is also about ethics, law, society, etc.

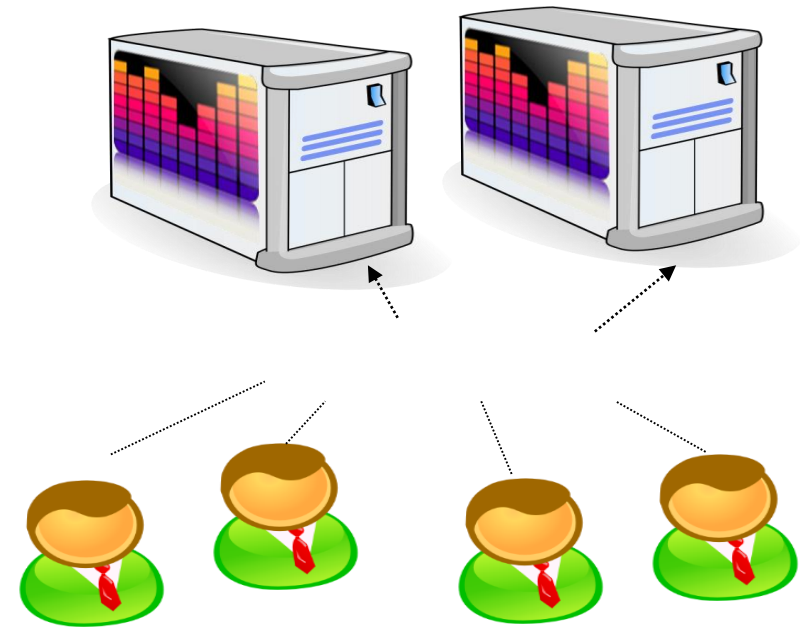
Provisioning and Allocation of Resources

Q: How do service creators get resources?

Provisioning



Allocation




Provisioning and Allocation Policies

Q: How many policies exist?

Q: How to select a policy?


Provisioning




When? From where?
How many?
Which type?
etc.



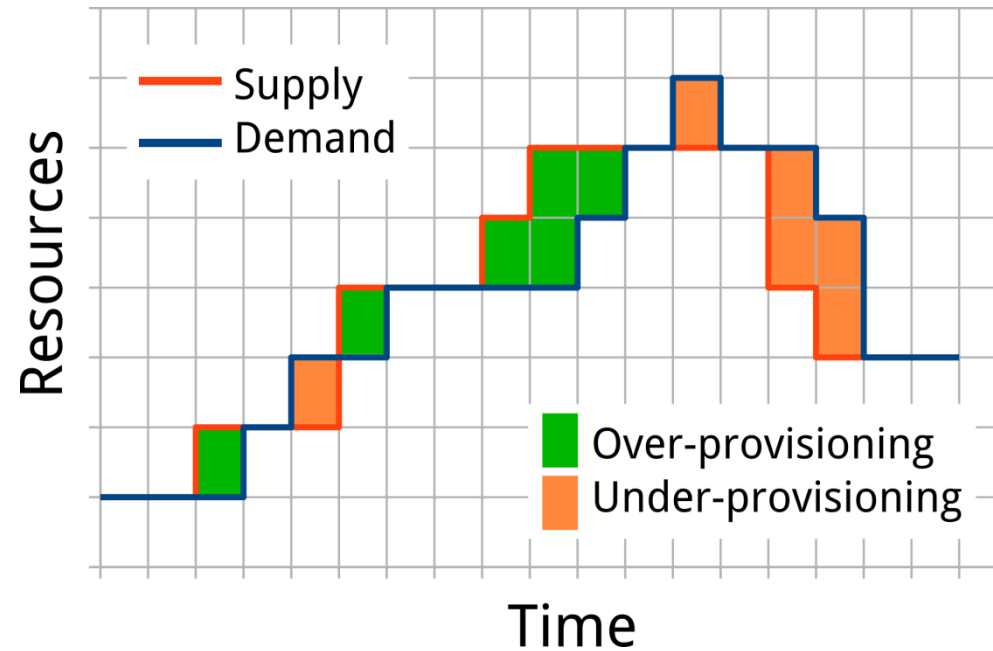
Allocation



When? Where?
etc.



Auto-Scalers = automatically provision resources, on-demand



Intro to Cloud Computing



5' — Pitch on Datacenter-Based Cloud Computing

5' — The Golden Age of Datacenters

20' — **A Delft View on Datacenters**

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- **The main challenges and techniques**

35' — **Making Clouds Tick**



- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10' — Reality Check

Here or @home

Interactive

Scientific Challenges to Get This Done



How to massivize?

- Super-scalable, super-flexible, yet efficient ICT infrastructure
- End-to-end automation of large-scale, simple and complex processes
- Dynamic, compute- and data-intensive workloads
- Evolving, heterogeneous hardware and software
- Strict performance, cost, energy, reliability, and fairness requirements
- ... all these, without needing much expertise from customers

↖ There's a lot we don't know
how to do yet...you can help!

The Scheduling Challenge



Cloud operator:

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

**Need scheduling policies for both
the cloud user and the cloud operator**

Cloud customer:


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

The Ecosystem Navigation Challenge

High-Level Language


Flume BigQuery SQL Meteor JAQL Hive Pig Sawzall Scope DryadLINQ AQL

**Need to support real users who choose their tools:
batch, workflows, stream, transactions, ...**




The "Big Cake" Challenge In the Datacenter

Online Social Networks



= Hadoop / MapReduce framework

Financial Analysts




Need multi-tenant, self-aware schedulers and resource managers

Universe Explorers



Multiple frameworks = Isolation, especially performance

Big Data Enthusiast



Jevons Effect: More Efficient, Less Capable?

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

**Need to be more efficient in how we use our resources,
(also educate others to not abuse “infinite” capacity)**

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, ...

Intro to Cloud Computing



5' — Pitch on Datacenter-Based Cloud Computing

5' — The Golden Age of Datacenters

20' — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35' — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10' — Reality Check

Here or @home

Interactive

Intro to Cloud Computing



5' — Pitch on Datacenter-Based Cloud Computing

5' — The Golden Age of Datacenters

20' — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35' — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10' — Reality Check

Here or @home

Interactive

The Scheduling Challenge



Cloud operator:

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

**Need scheduling policies for both
the cloud user and the cloud operator**

Cloud customer:

**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



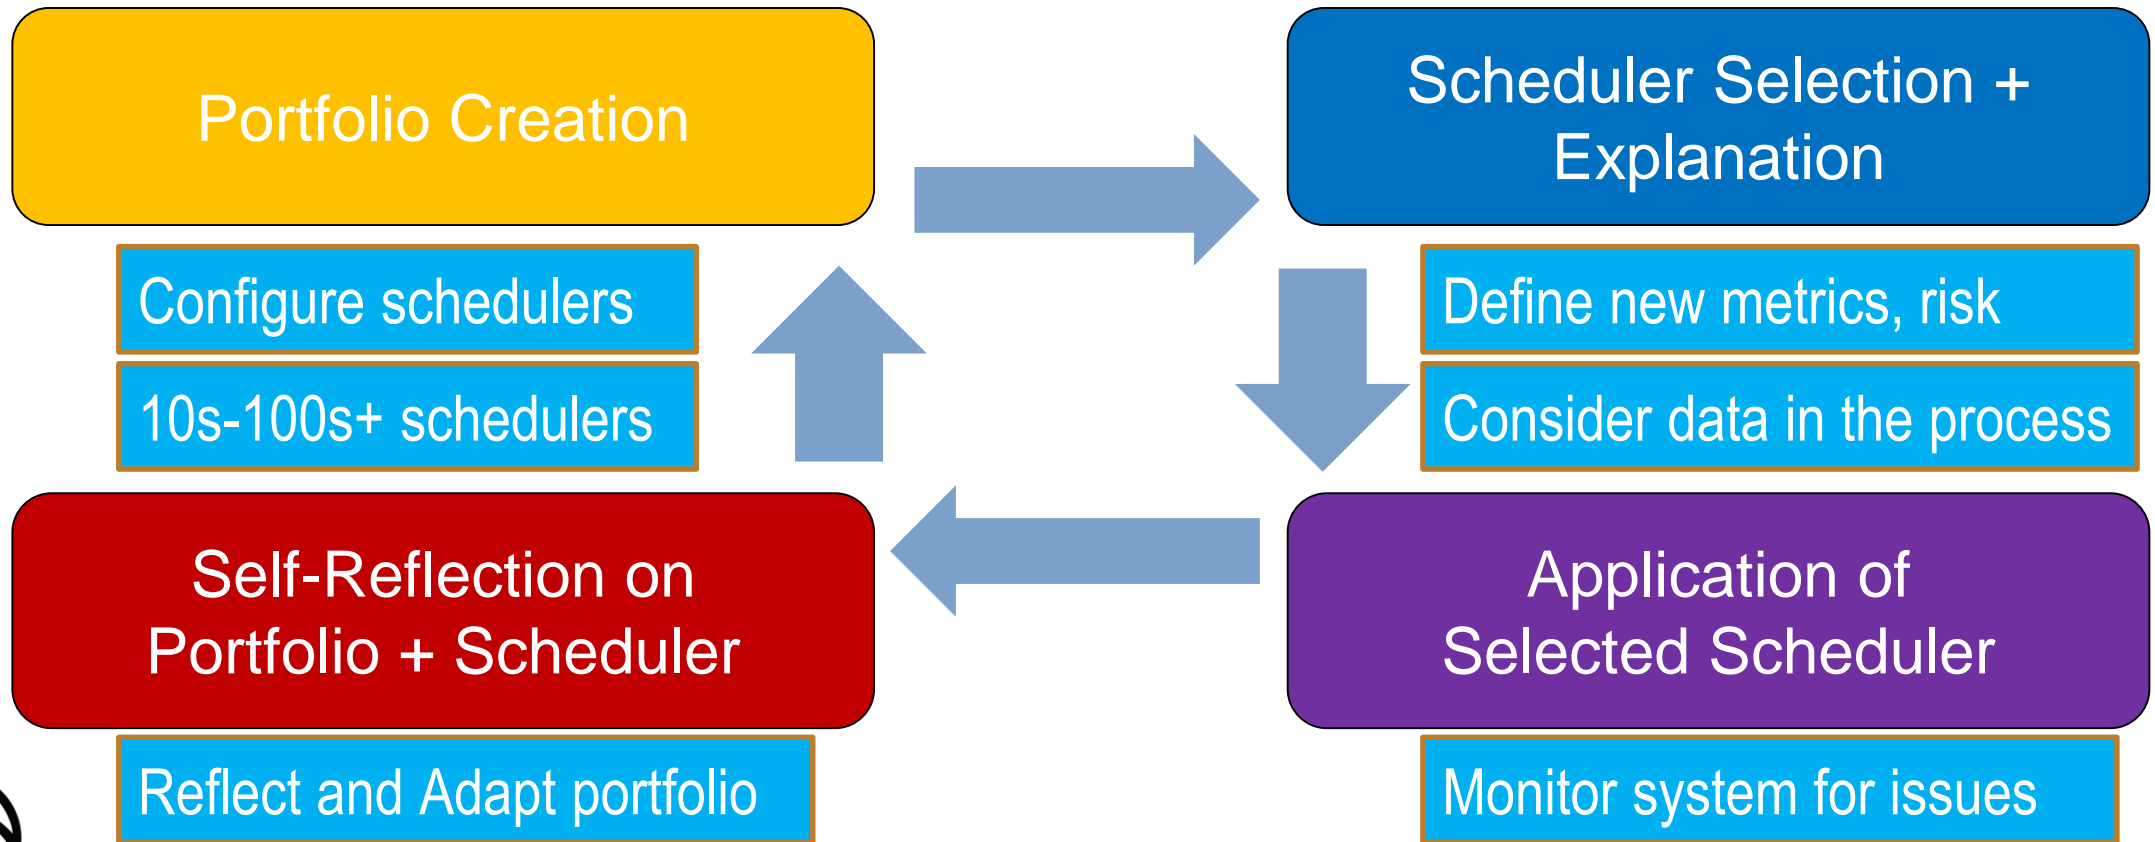
1. Create a set of schedulers (resource provisioning and allocation policies)
2. Select active scheduler online, apply for the next period, analyze results

(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

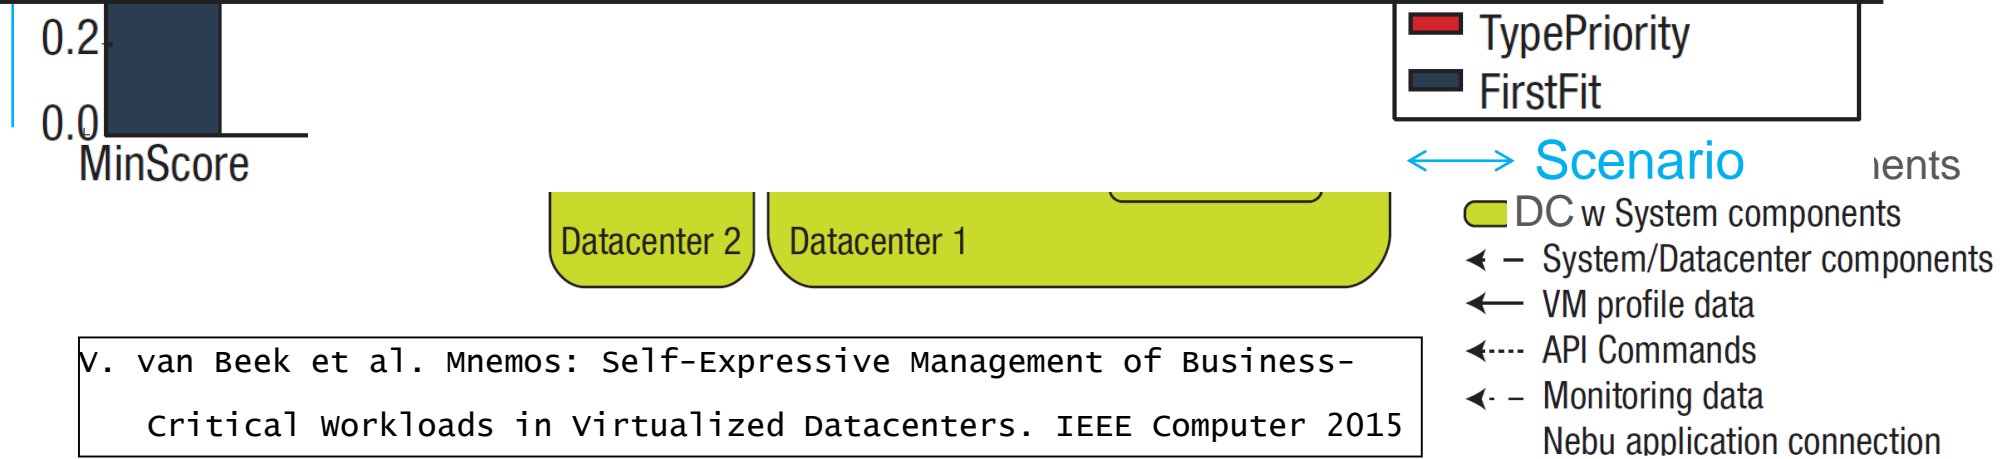


Portfolio Scheduling in Practice: Massive Datacenters



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. Memos: Self-Expressive Management of Business-critical workloads in virtualized Datacenters. IEEE Computer 2015



Intro to Cloud Computing



5' — Pitch on Datacenter-Based Cloud Computing

5' — The Golden Age of Datacenters

20' — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35' — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10' — Reality Check

Here or @home

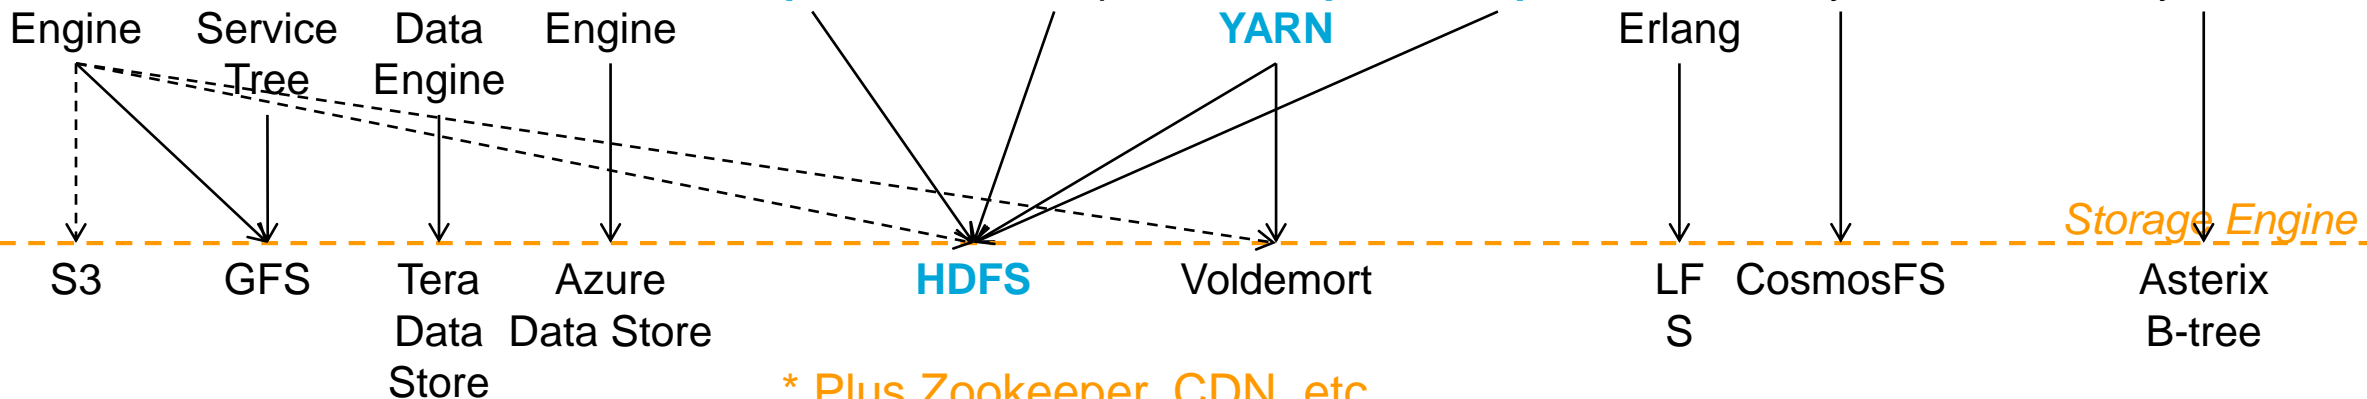
Interactive

The Ecosystem Navigation Challenge

High-Level Language

Flume BigQuery SQL Meteor JAQL Hive Pig Sawzall Scope Dryad LINQ AQL

**Need to support real users who choose their tools:
batch, workflows, stream, transactions, ...**





David Villegas
FIU/IBM



Athanasios
Antoniou



Alexandru
Iosup



Dick
Epema

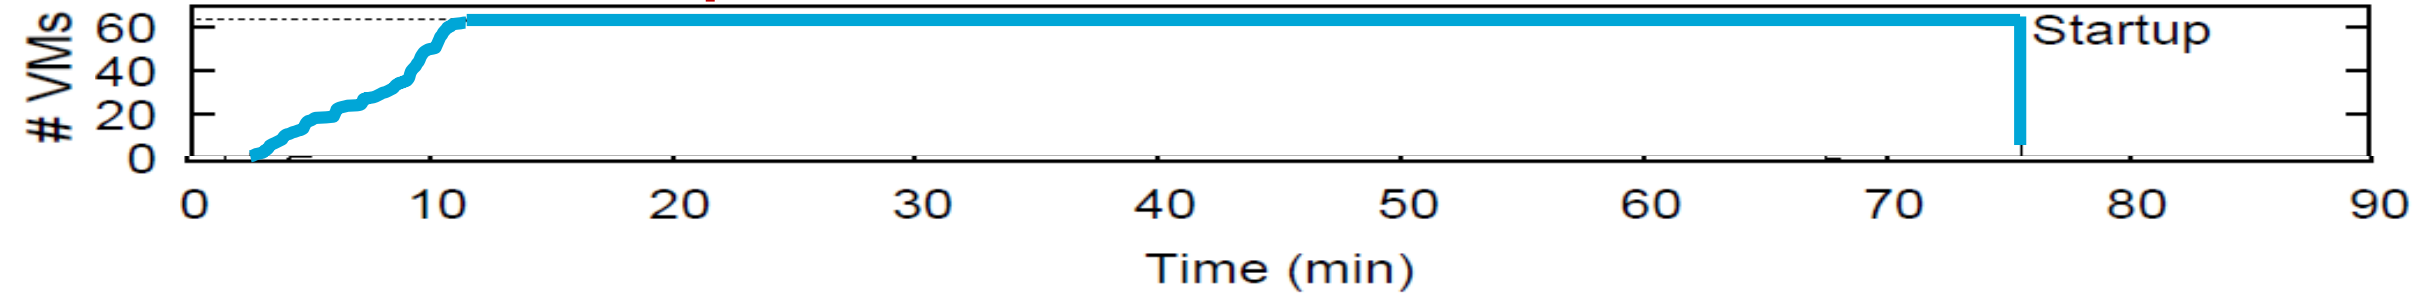
IaaS Provisioning and Allocation

Design of new policies and real-world experiments to compare with alternatives


Villegas, Antoniou, Sadjadi, Iosup. An Analysis of Provisioning and Allocation Policies for Infrastructure-as-a-Service Clouds, CCGrid 2012.

Use Case: Provisioning Policies, Compared

Startup



OnDemand




Requested Instances -----

Accessible Instances —————

Use Case: Provisioning Policies, Compared Environments (values for 2012 study, 2016 study much larger)

System	Hardware	VIM	Hypervisor	Max VMs
DAS4/Delft	20 Dual quad-core 2.4 GHz 24 GB RAM 2x1 TB storage	OpenNebula	KVM	64
FIU	7 Pentium 4 3.0 GHz 5 GB RAM 340 GB Storage	OpenNebula	Xen™	7
Amazon EC2	unkown/various	-	Xen™	20

Use Case: Two Provisioning Policies, Compared Workloads



Use Case: Two Provisioning Policies, Compared

Metrics for comparison

- Job Slowdown (**JSD**): Ratio of actual runtime in the cloud and the runtime in a dedicated non-virtualized environment

- Charged Cost (**C_c**)

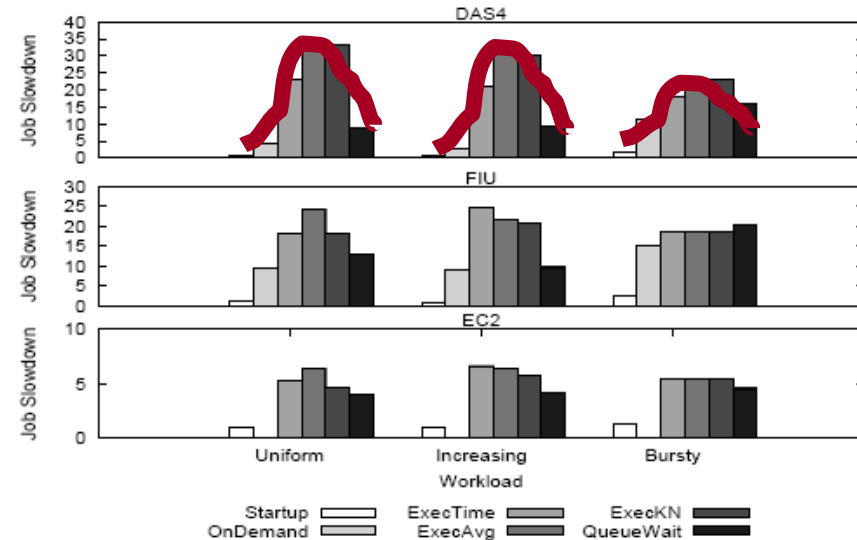

Q: Charged cost vs Total RunTime?

$$C_c(W) = \sum_{i \in \text{leased VMs}} [t_{stop}(i) - t_{start}(i)]$$

- Utility (**U**)

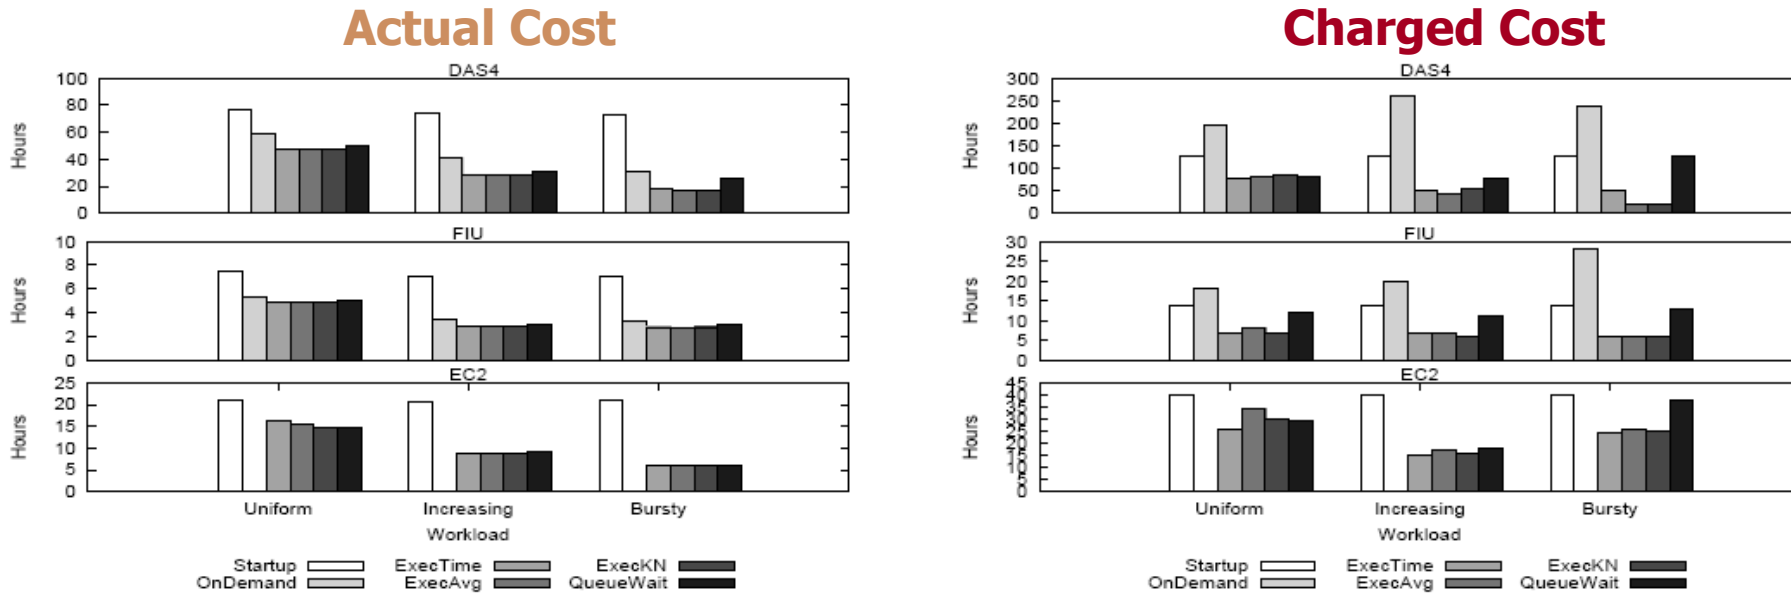
$$U(W) = \frac{SU_1(W)}{C_c(W)}$$

Performance Metrics





- Makespan very similar
- Very different job slowdown

Cost Metrics



- Very different results between actual and charged
 - Cloud charging function an important selection criterion
- All policies better than Startup in actual cost
- Policies much better/worse than Startup in charged cost

Compound Metrics



- Trade-off Utility-Cost still needs investigation
- **Performance or Cost, not both:**
the policies we have studied improve one, but not both



Alexey
Ilyushkin



Ahmed
Ali-Eldin



Nikolas
Herbst



Alessandro
Papadopoulos



Bogdan
Ghiț



Dick
Epema



Alexandru
Iosup

Auto-Scaling

Experimental Performance Evaluation of Autoscaling Policies for Complex Workflows

Best Paper Candidate

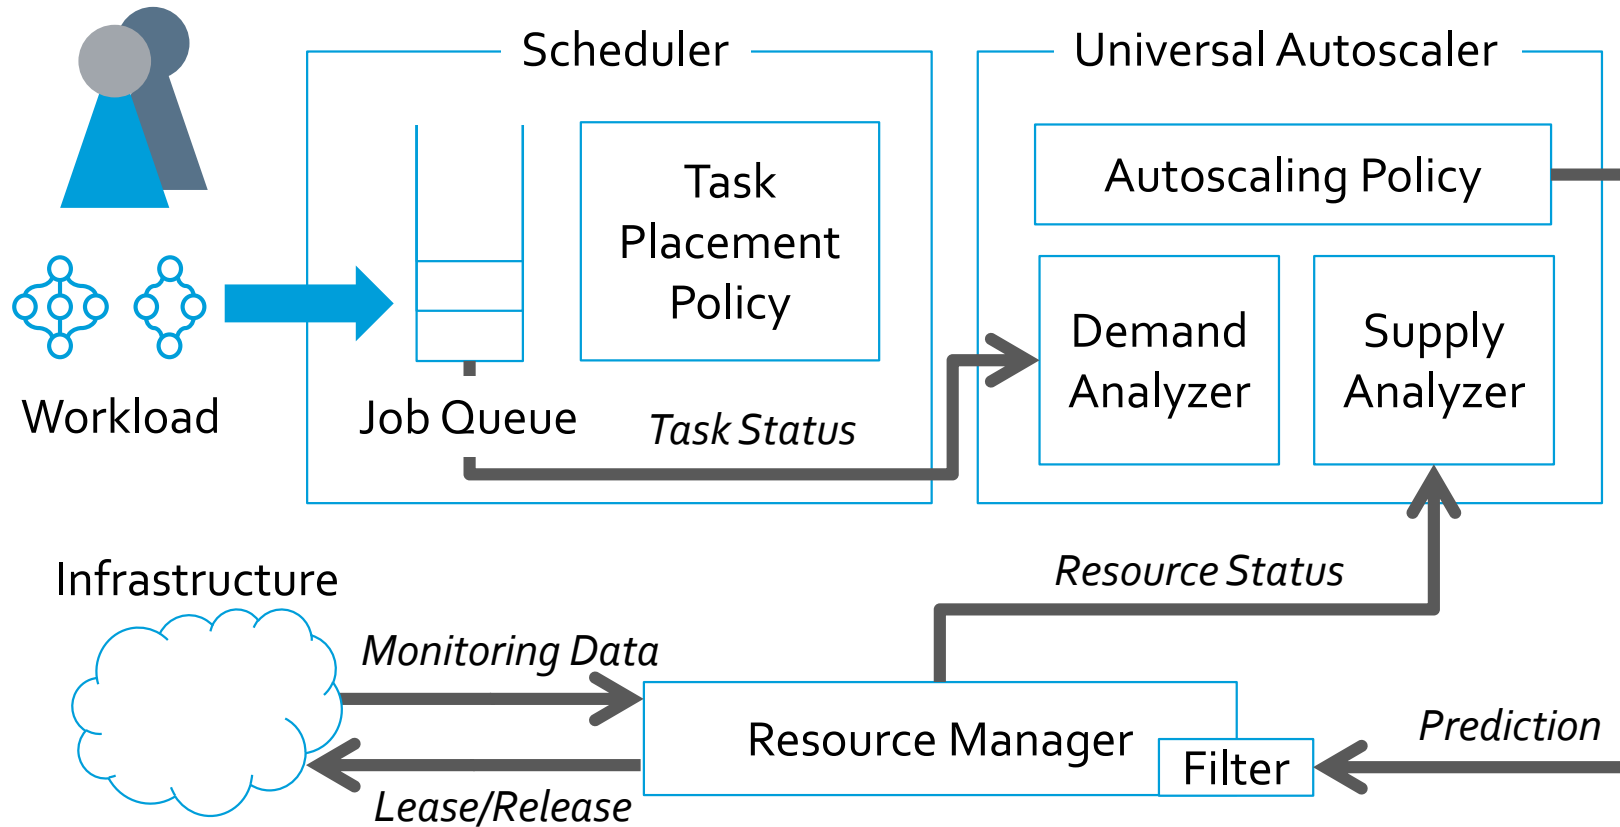
Ilyushkin, Ali-Eldin, Herbst, Papadopoulos, Ghiț, Epema, Iosup. An Experimental Performance Evaluation of Autoscaling Policies for complex workflows. ICPE 2017

Our Approach

A comprehensive method for evaluating and comparing autoscalers

- A **model** for elastic cloud platform
- A set of relevant **metrics** for assessing autoscaler performance
- A set of general and workflow-specific **autoscalers**
- Three **comparison methods** for autoscalers
- **Real experiments** with up to 50 VMs in OpenNebula on DAS supercomputer

Elastic Cloud Platform



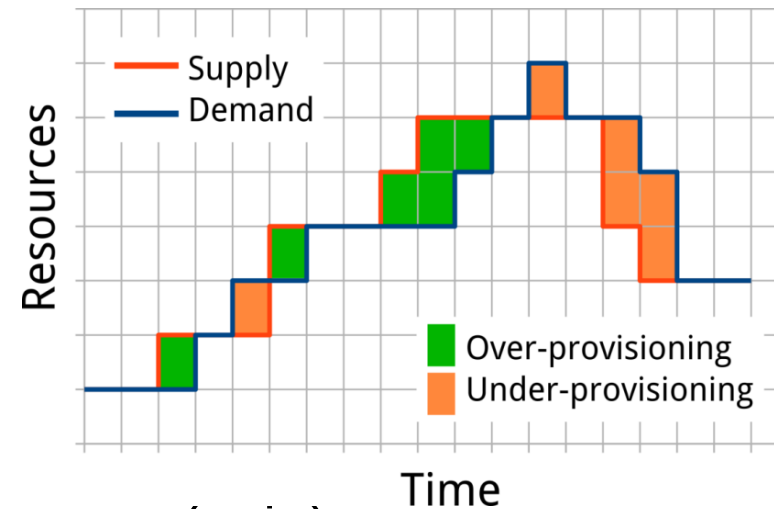
Performance Metrics

System-oriented elasticity metrics

- Accuracy (also normalized by actual demand)
- Wrong-Provisioning Timeshare
- Instability


User-oriented metrics

- Elastic Slowdown
- Average Number of Utilized Resources (gain)



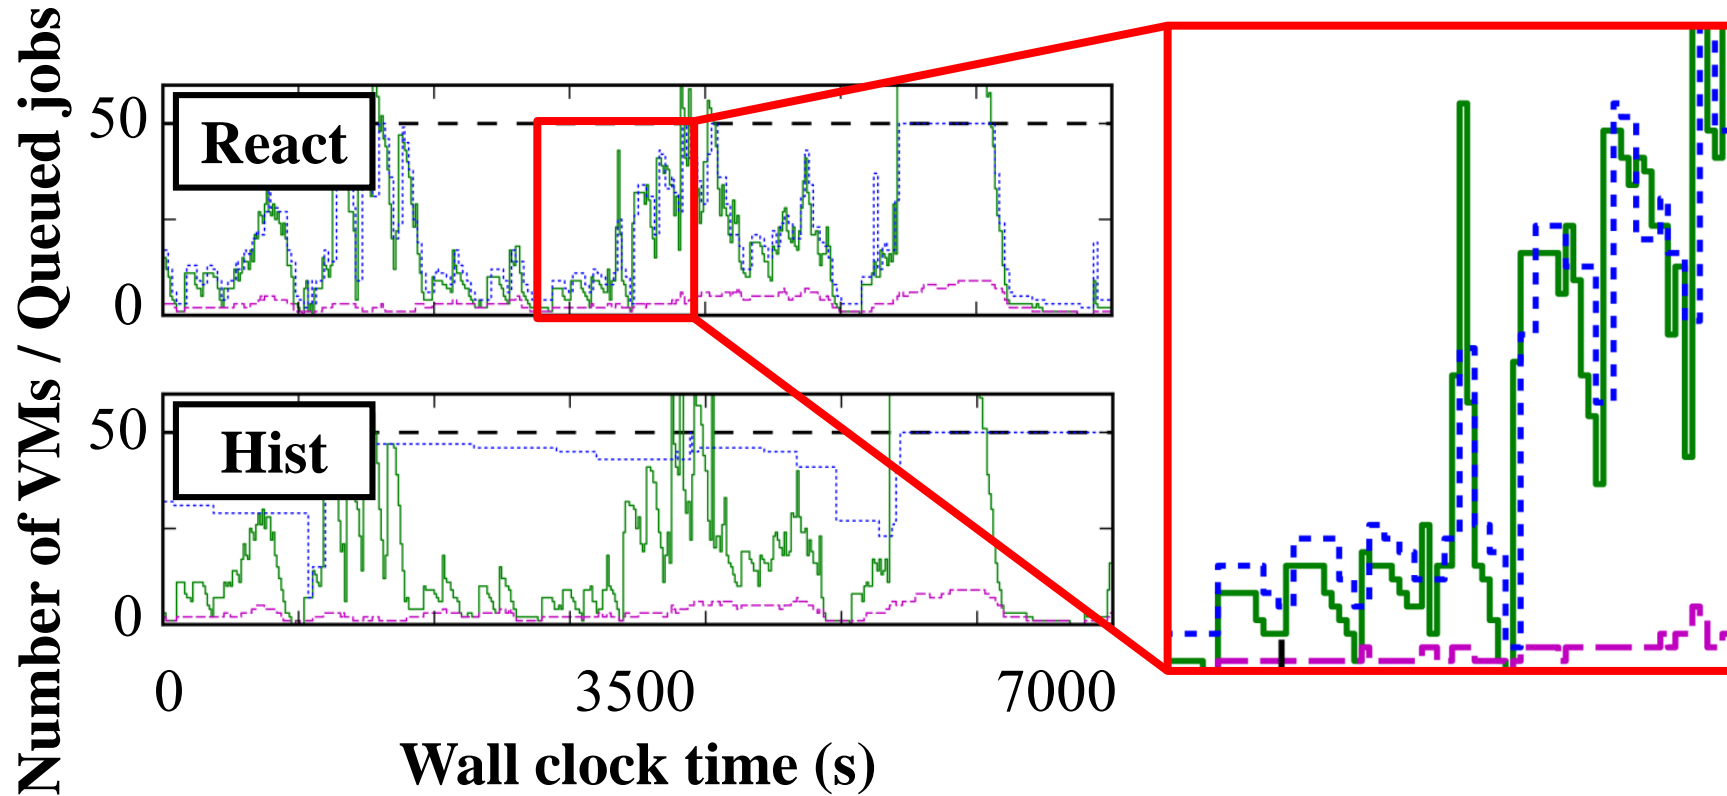
Autoscaling Policies

Timeliness of the Information



Information Source

Experimental Results




— Demand (VMs) - - - - Supply (idle and busy VMs)
- - - - Queue length (workflows)

Which Policy is the Best?

Methods for aggregation of metrics

- **Pairwise Comparison** – pairwise compare metrics between autoscalers
- **Fractional Difference Comparison** – compare autoscalers with an ideal case based on the experimental results
- **Aggregated System-oriented Elasticity and User Metrics**
(by Fleming et al.)
Compute speedup ratios and then average the speedups using an unweighted geometric mean

Which Policy is the Best?



The horizontal scale is cropped!

Conclusion

1. We developed a method to compare different autoscalers
2. General autoscalers can achieve similar performance as workflow-specific autoscalers (surprising)
3. No autoscaler is the best:
Our workflow-specific Plan autoscaler wins 4 out of 5 competitions but is not the best overall
4. The correct choice of an autoscaler is important but significantly depends on the application type
5. Correct parameterization of general autoscalers is very important

Intro to Cloud Computing



5' — Pitch on Datacenter-Based Cloud Computing

5' — The Golden Age of Datacenters

20' — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35' — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge


10' — Reality Check

Here or @home

Interactive


The "Big Cake" Challenge In the Datacenter

Online Social Networks



= Hadoop / MapReduce framework

Financial Analysts




Need multi-tenant, self-aware schedulers and resource managers

Universe Explorers





Multiple frameworks = Isolation, especially performance

Big Data Enthusiast



Dynamic Big Data Processing

Fawkes = Elastic MapReduce



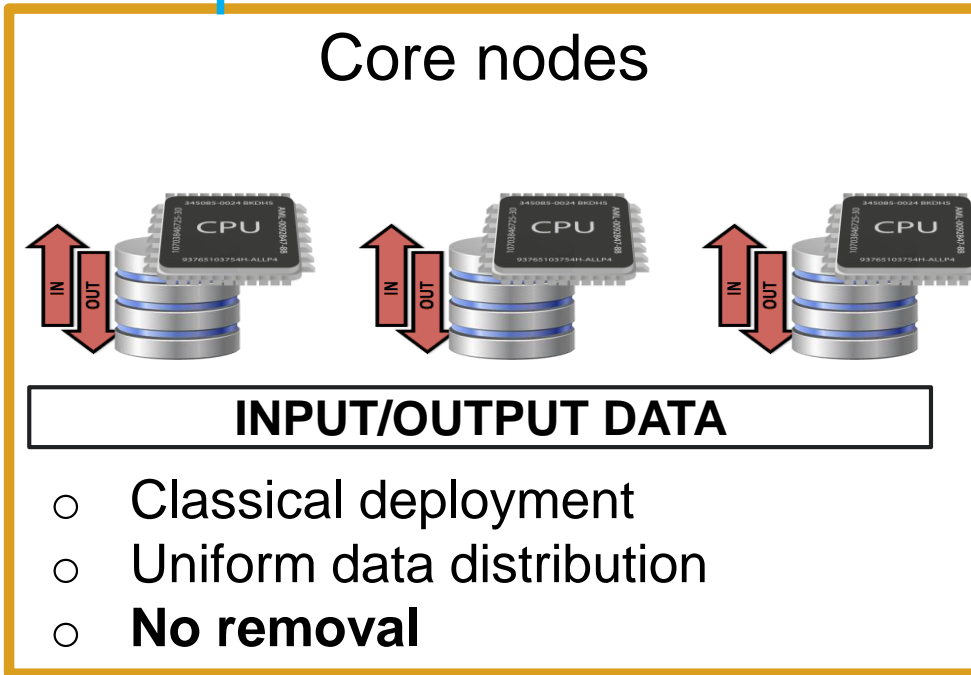
Job submissions

Frameworks

Resource manager

Infrastructure

Elasticity for MapReduce Frameworks



Transient nodes (TR)



NO DATA

- No local storage
- R/W from/to core nodes
- **Instant removal**

Trans-core nodes (TC)





OUTPUT DATA

- Local storage, no input
- Only R from core nodes
- **Delayed removal**

Fawkes in a Nutshell [1/2]

Because workloads may be time-varying:

- Poor resource utilization
- Imbalanced service levels



1. Fair framework size:

$$s_i = \frac{w_i}{w_1 + w_2 + w_3}, \quad i = 1, 2, 3$$

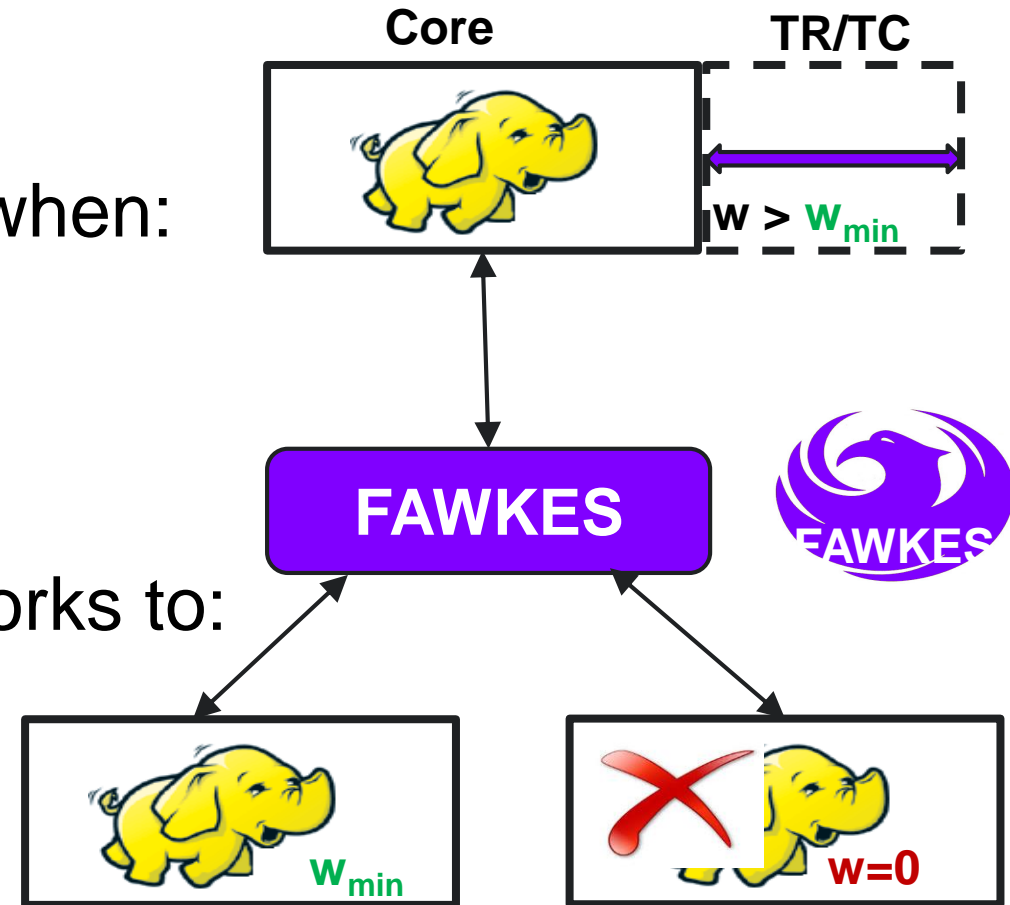
Fawkes in a Nutshell [2/2]

2. Updates dynamic weights when:

- New frameworks arrive
- Framework states change

3. Shrinks and grows frameworks to:

- Allocate **new** frameworks
- Give fair shares to existing frameworks
- **Eliminate unused** frameworks



MapReduce Applications Tested with Fawkes

Application	Type	Input	Output
Wordcount (WC)	CPU	200 GB	5.5 MB
Sort (ST)	Disk	200 GB	200 GB
PageRank (PR)	CPU	50 GB	1.5 MB
K-Means (KM)	Both	70 GB	72 GB
TrackerOverTime (TT)	CPU	100 GB	3.9 MB
ActiveHashes (AH)	Both	100 GB	90 KB
BTWorld (BT)	Both	100 GB	73 GB



Synthetic benchmarks:

- HiBench suite
- Single applications
- Random datasets

Real-world applications:

- BTWorld workflow
- 14 Pig queries
- BitTorrent monitoring data

Performance of dynamic MapR

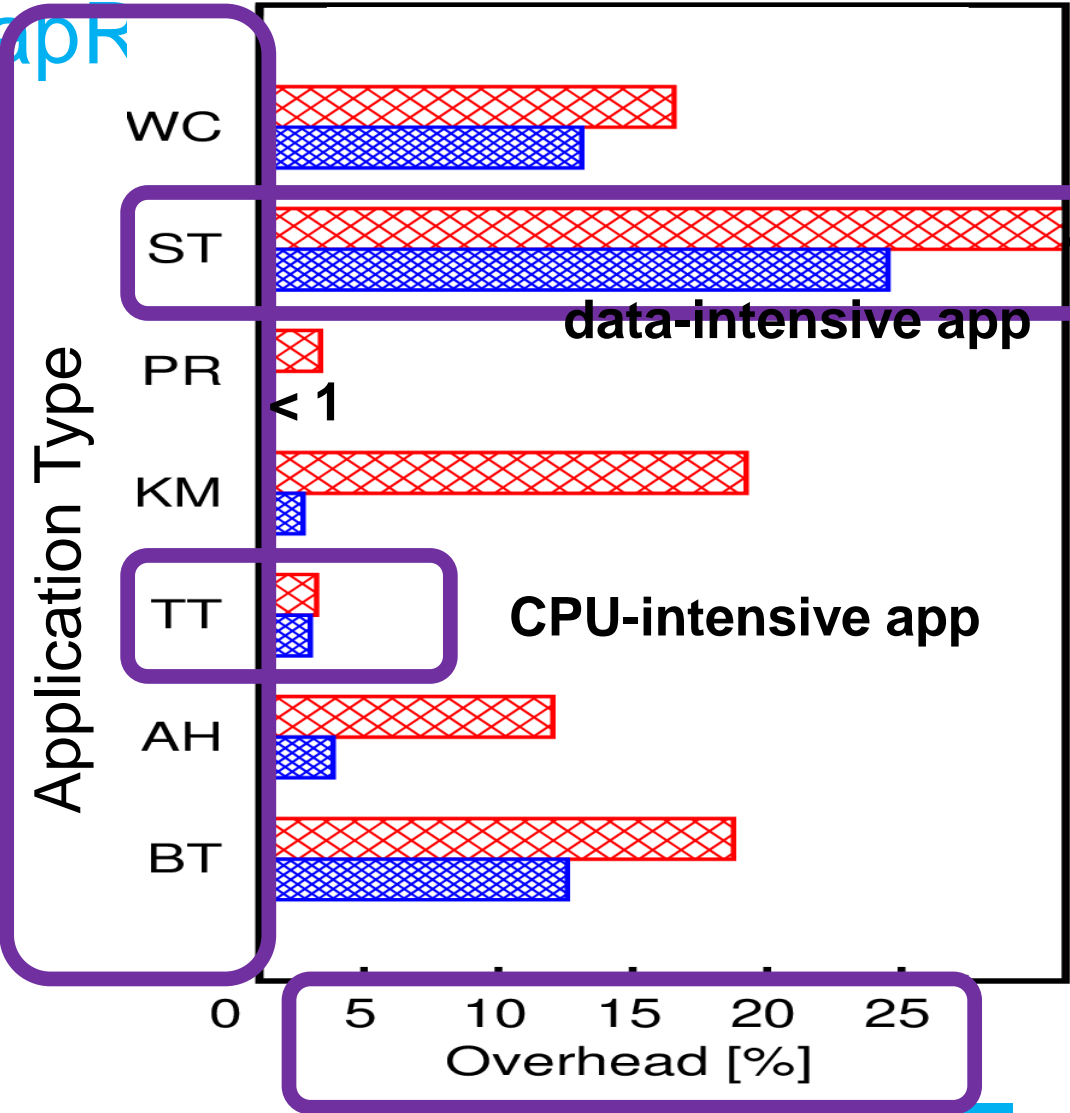
10 core + 10xTR 
10 core + 10xTC 
vs.
20 core nodes (baseline)

TRansient - good for compute-intensive workloads.

TCore - needed for disk-intensive workloads.

Dynamic MapReduce:
< 25% overhead

Fawkes also reduces imbalance



Intro to Cloud Computing



5' — Pitch on Datacenter-Based Cloud Computing

5' — The Golden Age of Datacenters

20' — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35' — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10' — Reality Check

Here or @home

Interactive

Jevons Effect: More Efficient, Less Capable

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

**Need to be more efficient in how we use our resources,
(also educate others to not abuse “infinite” capacity)**

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, ...

Existing Graph-Processing Systems: *Either Distributed or Heterogeneous*

- **Distributed CPU-based** systems cannot use additional computational power of accelerators



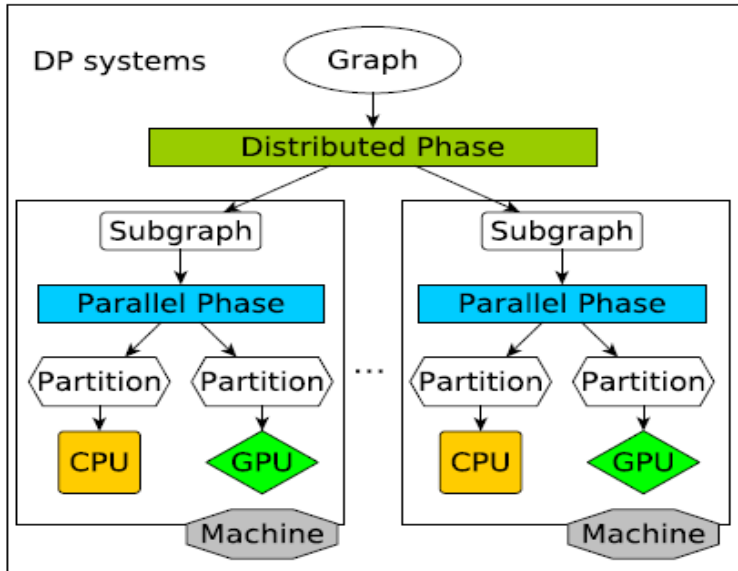
- **GPU-enabled** systems are (mostly) single-machine systems, cannot handle large-s



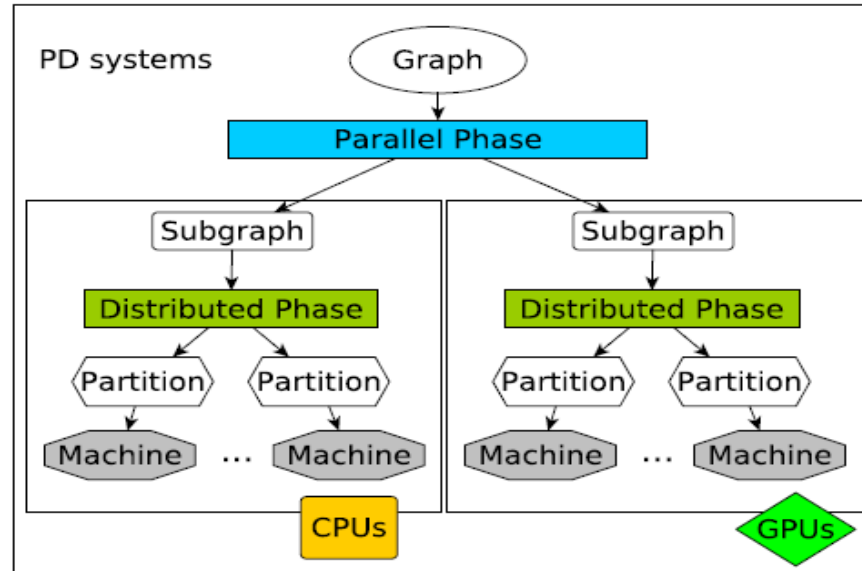
Y. Guo, A. L. Varbanescu, D. Epema, and A. Iosup, "Design and Experimental Evaluation of Distributed Heterogeneous Graph-Processing Systems," Submission to *CCGrid*, 2016.

© 2017 Alexandru Iosup. All rights reserved.

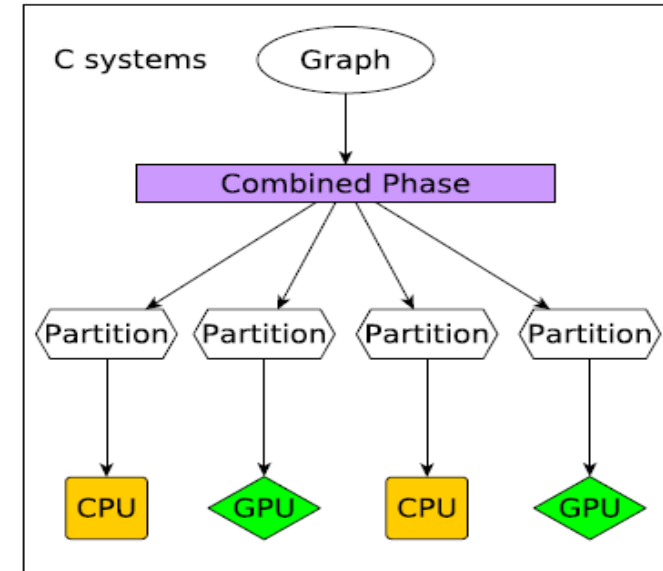
Our approach: 3 Families of Distributed *and* Heterogeneous (CPU+GPU) Graph-Processing Systems



Distributed-then-Parallel (DP) Systems

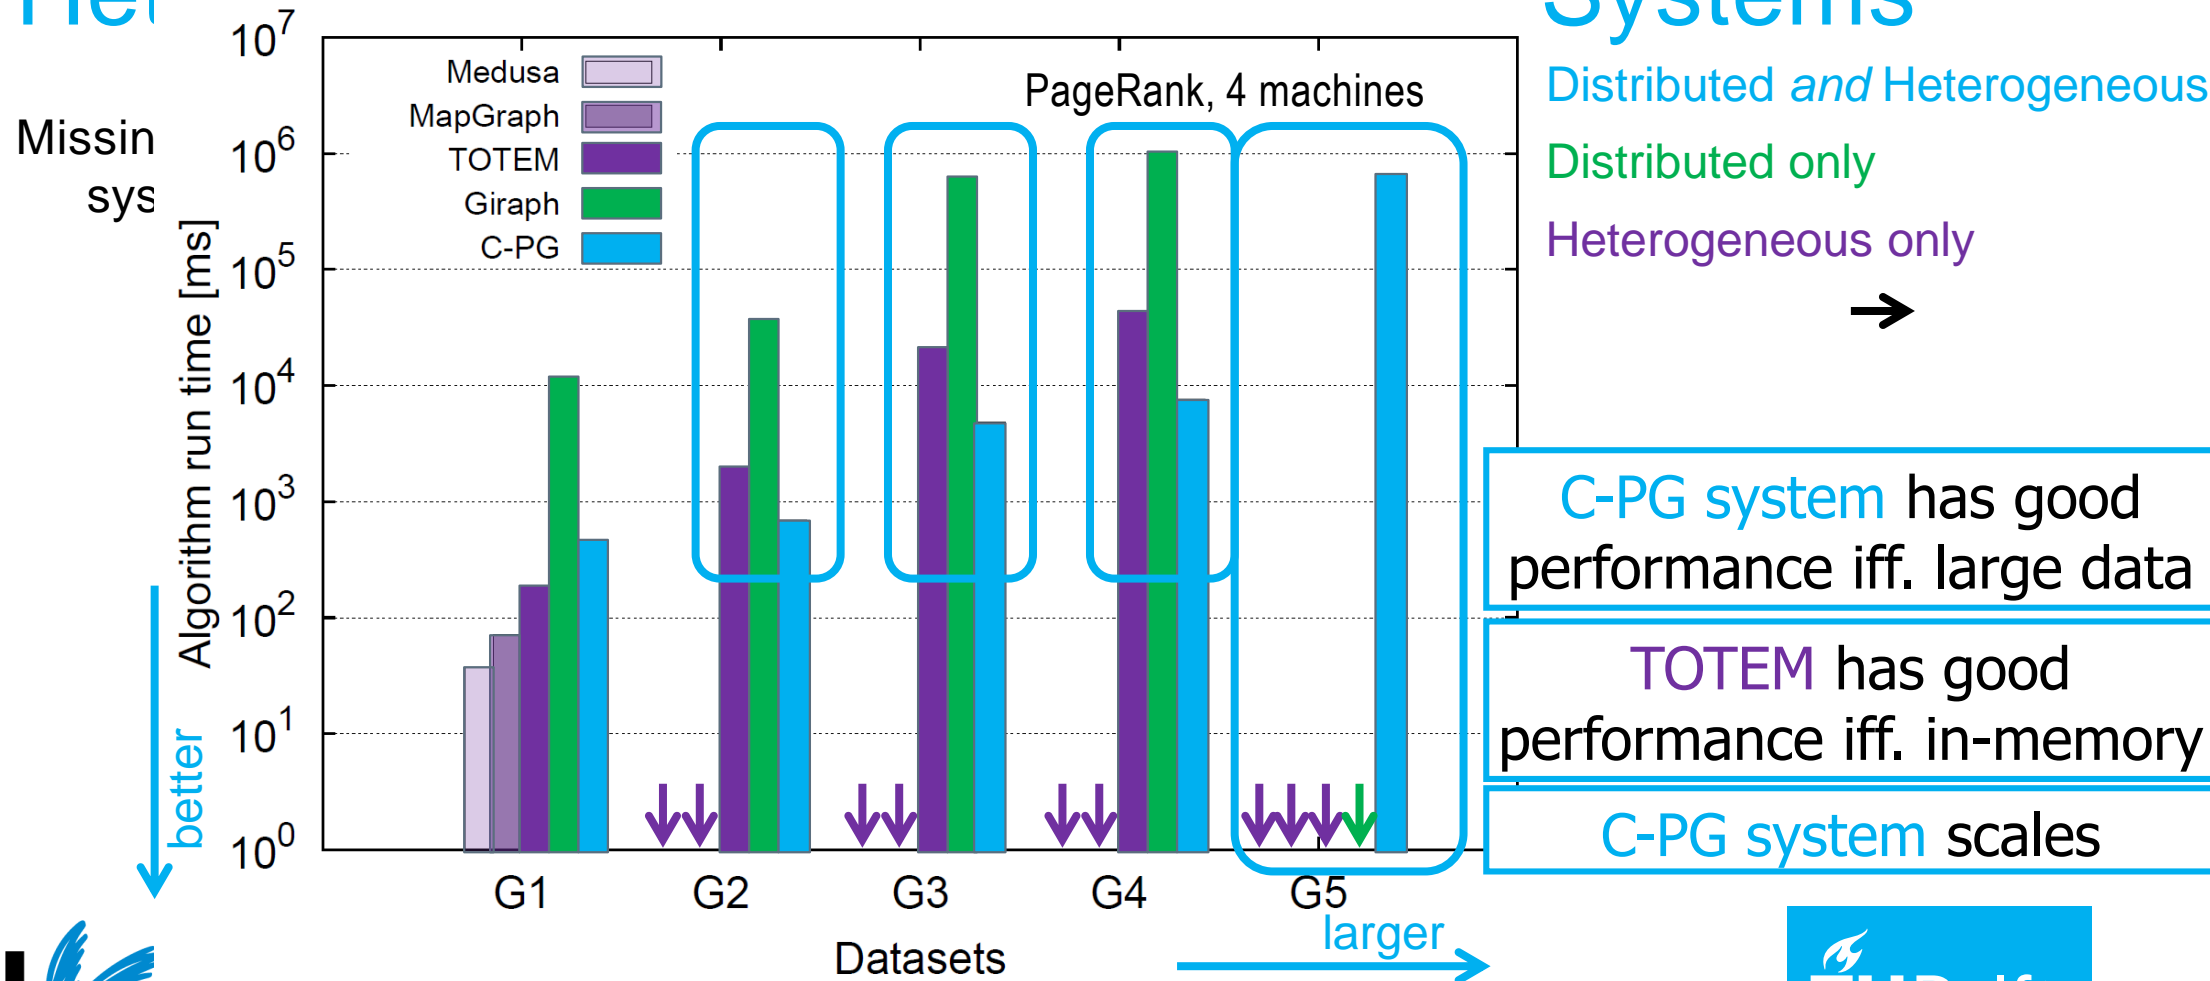


Parallel-then-Distributed (PD) Systems



(Combined Par.-and-Distributed (C) Systems

Promising Results for Distributed *and* Heterogeneous Graph Processing Systems



Intro to Cloud Computing



5' — Pitch on Datacenter-Based Cloud Computing

5' — The Golden Age of Datacenters

20' — A Delft View on Datacenters

- The core idea of datacenter computing
- The main enabling technologies for datacenter computing
- The main challenges and techniques

35' — Making Clouds Tick

- Addressing the Scheduling challenge
- Addressing the Ecosystem Navigation challenge
- Addressing the Big Cake challenge
- Addressing the Efficiency challenge

10' — Reality Check

Here or @home

Interactive Masterclass

The Real IaaS Cloud

“The path to abundance”

On-demand


Cheap

Great fit



<http://www.flickr.com/photos/dimitrisotiropoulos/4204766418/>


VS



Tropical Cyclone Nargis (NASA, ISSS, 04/29/08)



- “The killer cyclone”
- **Not so great performance for scientific applications** (compute- or data-intensive)

Animoto: Video App on Amazon EC2



Zynga zCloud: Hybrid Self-Hosted/EC2

- After Zynga had large scale
- More efficient self-hosted servers



Other Cloud Customers



- 218 virtual CPUs
- 9TB/2TB block/S3 storage
- 6.5TB/2TB I/O per month



Customers in 190 Countries



Our Industry Collaborators



AZAVISTA



Microsoft



Take-Home Message


The Golden Age of Datacenters

Cloud computing = lease vs self-own

- On-Demand, Pay-per-Use, Elastic, Pooled, Automated, QoS
- Owner vs User perspective

Important New Challenges

1. The scheduling challenge
2. The ecosystem navigation challenge
3. The big cake & 4. Jevons (Efficiency) challenges



Recommended Reading <http://atlarge.science>

Elastic Big Data and Computing

- van Beek et al.: Self-Expressive Management of Business-Critical Workloads in Virtualized Datacenters. IEEE Computer 48(7): 46-54 (2015)
- B. Ghit, N. Yigitbasi (Intel Research Labs, Portland), A. Iosup, and D. Epema. Balanced Resource Allocations Across Multiple Dynamic MapReduce Clusters. SIGMETRICS 2014
- L. Fei, B. Ghit, A. Iosup, D. H. J. Epema: KOALA-C: A task allocator for integrated multicluster and multicloud environments. CLUSTER 2014: 57-65
- K. Deng, J. Song, K. Ren, A. Iosup: Exploring portfolio scheduling for long-term execution of scientific workloads in IaaS clouds. SC 2013: 55
- Iosup et al.: On the Performance Variability of Production Cloud Services. CCGRID 2011: 104-113
- Iosup et al.: Performance Analysis of Cloud Computing Services for Many-Tasks Scientific Computing. IEEE Trans. Parallel Distrib. Syst. 22(6): 931-945 (2011)

Disclaimer: images used in this presentation obtained via Google Images.

- Images used in this lecture courtesy to many anonymous contributors to Google Images, and to Google Image Search.
- Many thanks!