

Cloud Computing: Open Research Questions for the Next Decade




Alexandru Iosup

Parallel and Distributed Systems Group
Delft University of Technology
The Netherlands

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



Source: <http://dilbert.com/strips/comic/1997-11-22/>

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

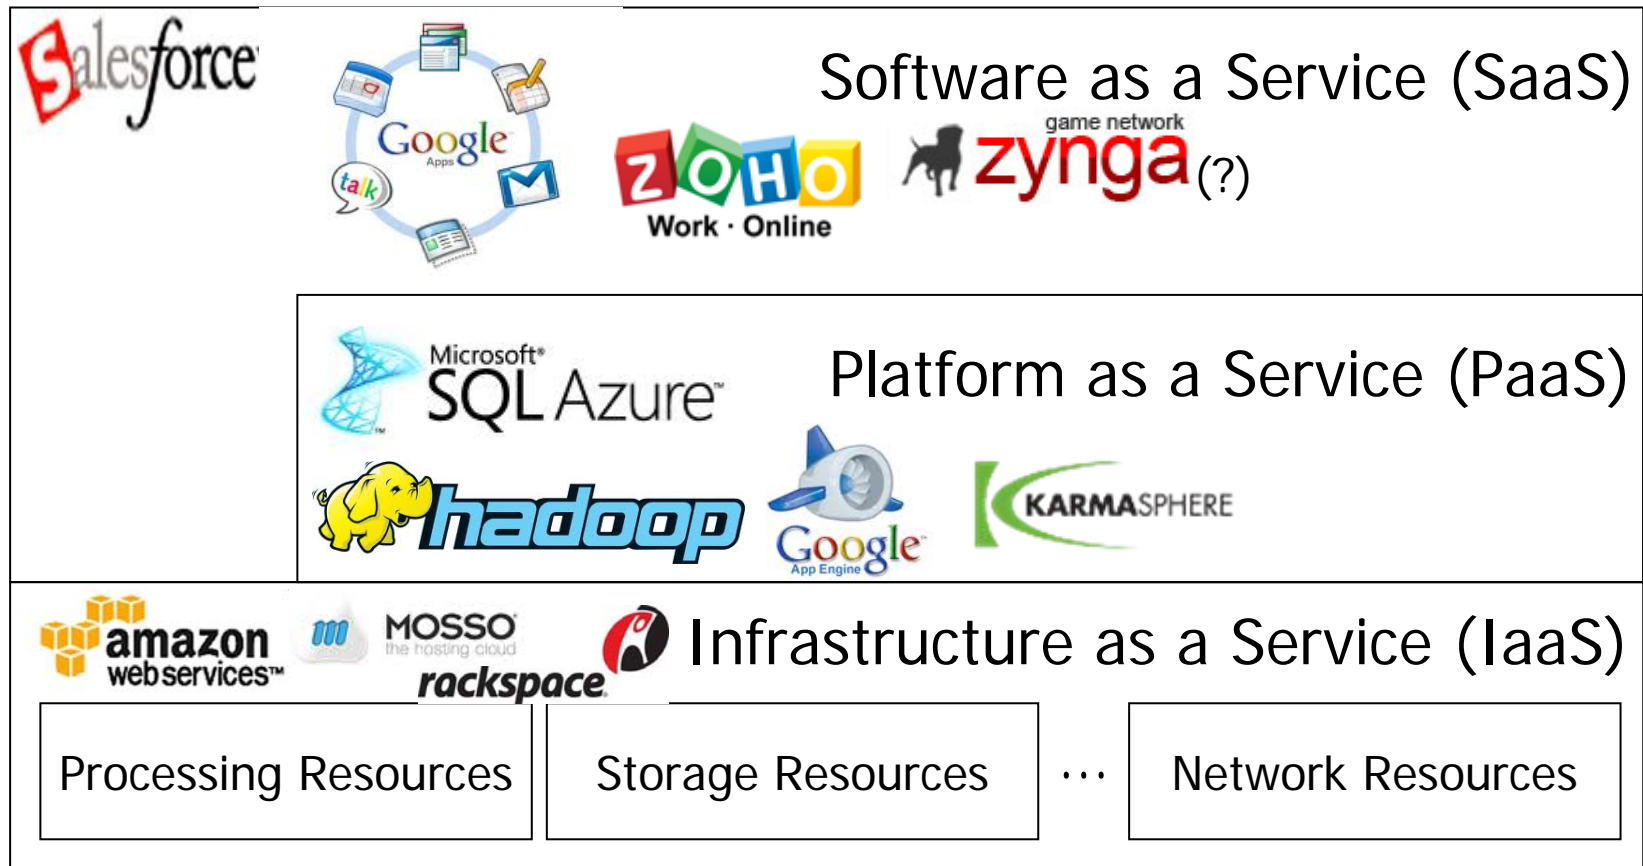
Virtualized HW + OS

MW = Middleware

What is Cloud Computing?

3. A Useful IT Service

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




What is Cloud Computing?


4. A Useful IT Deployment Model




Private cloud



Community cloud



Hybrid cloud



Public cloud

Legend

Serviced
Not Serviced

Source: Mell and Grance, NIST Spec.Pub. 800-145, Sep 2011.

What is Cloud Computing?

Time to Vote

1. A cloudy Buzzword
2. A Descendant of the Grid
3. A useful IT Service
4. A useful IT Deployment Model



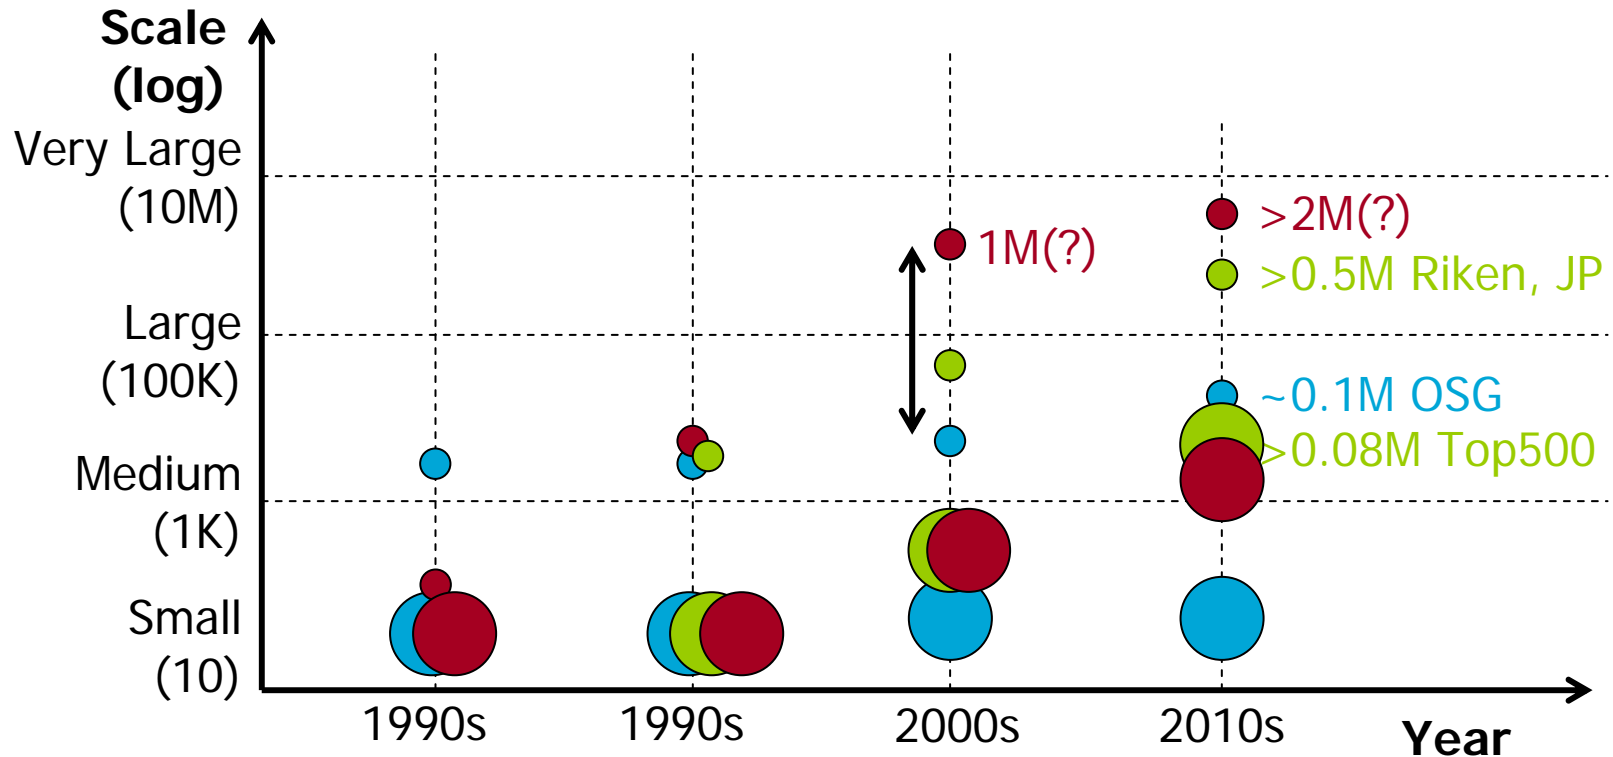
Other Votes

	QoS/ Elasticity/ Scalability	Cost/ ROI	Availability	Performance	Security	Data Control	Hybrid Cloud
Importance of Issue							
Contrail, thanks to Guillaume Pierre							
n=20: 5 academic, 5 IT services, 10 others	4.2/5	3.6/5		4.1/5	4.6/5	3.5/5	4.2/5
Concerns in Using Cloud Computing							
BITKOM'09/ CIO Survey/ Helmut Krcmar	12%	11%	25%	24%	45%	26%	26%

Agenda

1. Introduction to Cloud Computing
- 2. Is there a Role for Academics in Cloud Computing?**
3. Which Applications Need Cloud Computing?
4. How to Build a Toolbox for Cloud Computing Research?
5. How to Build the Core of Cloud Computing?
6. Conclusion

Is there a Role for Academics in Cloud Computing? [Other than as Users]



Legend

Industry (private clouds)

Academia (grids)

Government (supercomputers)

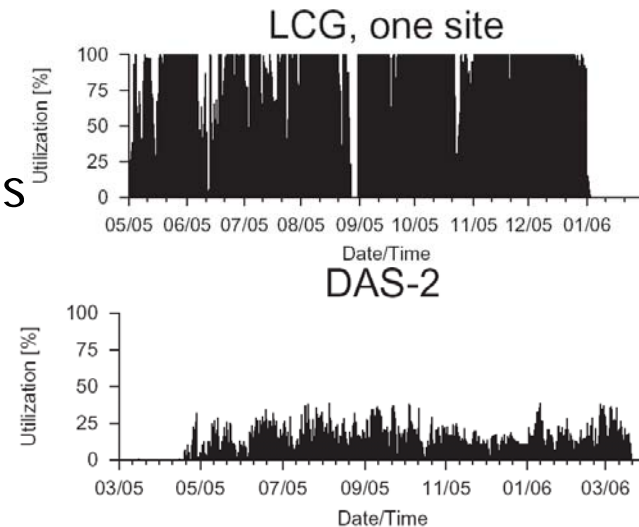
Circle size ~ instance count

Agenda

1. Introduction to Cloud Computing
2. Is there a Role for Academics in Cloud Computing?
- 3. Which Applications Need Cloud Computing?**
4. How to Build a Toolbox for Cloud Computing Research?
5. How to Build the Core of Cloud Computing?
6. Conclusion

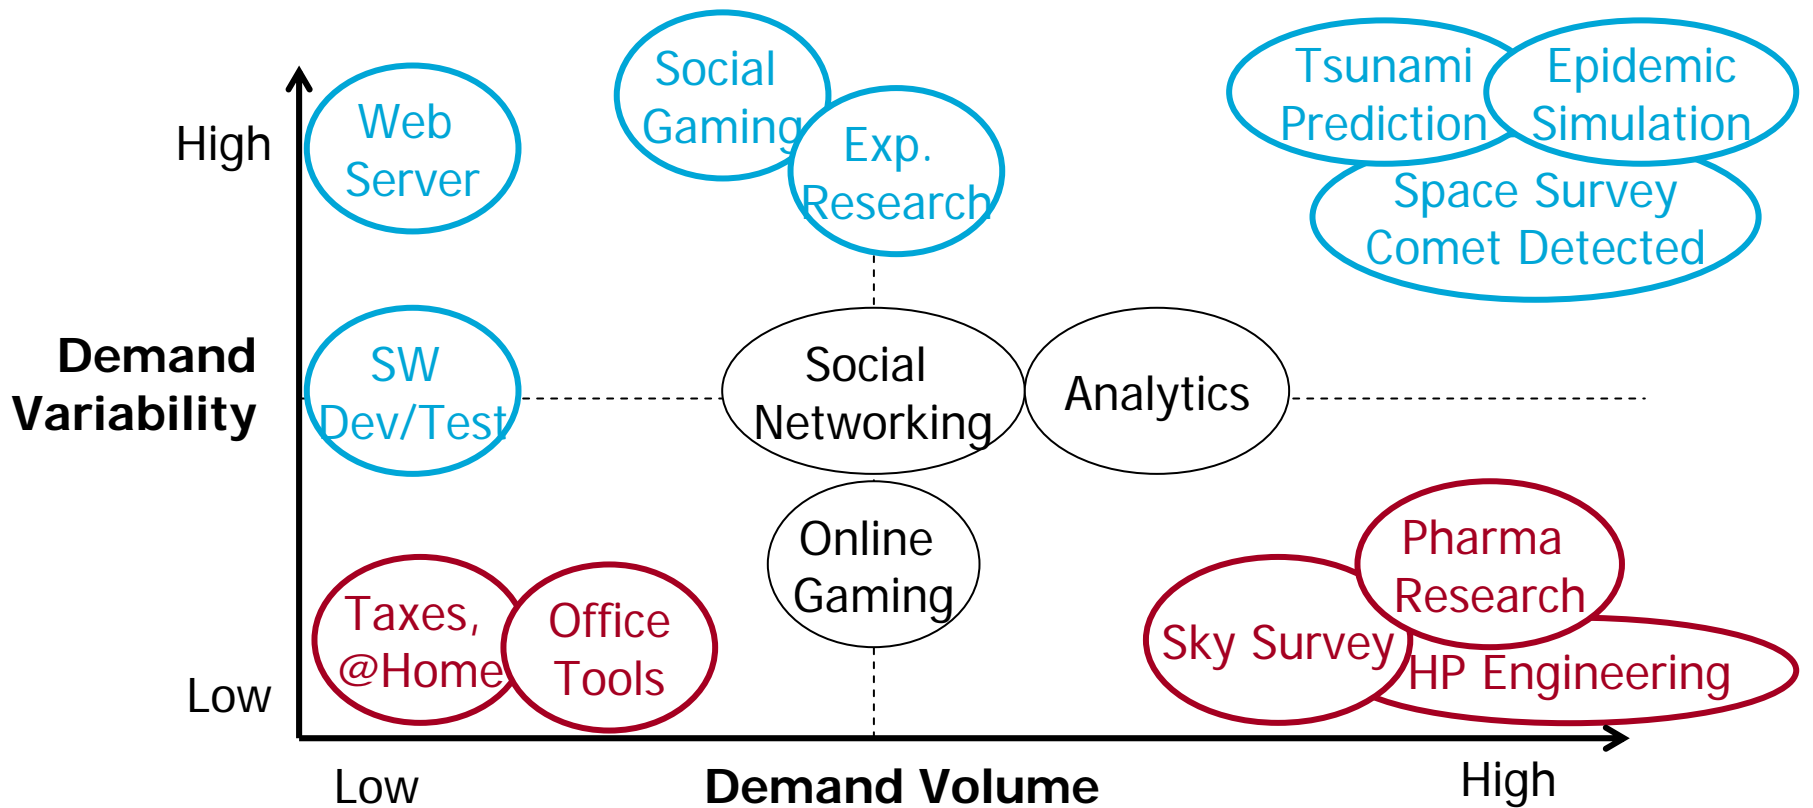
It Took 10 Years to Find that Workloads are different in Grids vs Parallel Envs.

- Traces: LCG, Grid3, TeraGrid, and DAS
 - long traces (6+ months), active environments (500+K jobs per trace, 100s of users), >4 million jobs
- Analysis
 - System-wide, VO, group, user characteristics
 - Environment, user evolution
 - System performance
- Selected findings
 - **Almost no parallel jobs!**



A. Iosup, C. Dumitrescu, D.H.J. Epema, H. Li, L. Wolters,
How are Real Grids Used? The Analysis of Four Grid Traces
and Its Implications, Grid 2006.

Which Applications Need Cloud Computing? A Simplistic View...




What's in a Name? MSG, MMOG, MMO, ...

250,000,000 active players
3BN hours/week world-wide

Massively Social Game =

online game with massive numbers of players (100K+), for which the **social interaction** improves the experience



1. Virtual world

Simulation,
content streaming
+

2. Content

Generation
+

3. Game data

Player activity, stats,
and (**social**)
relationships

Cloudfying Social Games



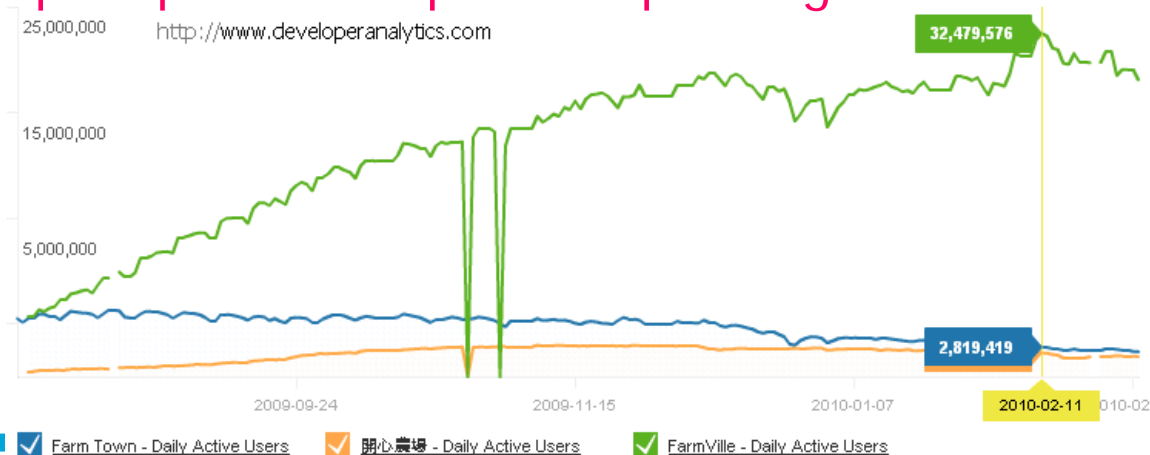
+



(Platform Challenge)

Build MSG platform that uses (mostly) cloud resources

- Close to players
- No upfront costs, no maintenance
- Compute platforms: multi-cores, GPUs, clusters, all-in-one!
- Performance guarantees
- Code for various compute platforms—platform profiling
- Misprediction=\$\$\$
- What services?
- Vendor lock-in?
- *My data*



Nae, Iosup, Prodan, Dynamic Resource Provisioning in Massively Multiplayer Online Games, IEEE TPDS, 2011.

Research Opportunity: Characterize and Model Cloud Apps


- Looking for invariants

- Wr [%] ~40% Total IO, but absolute values vary

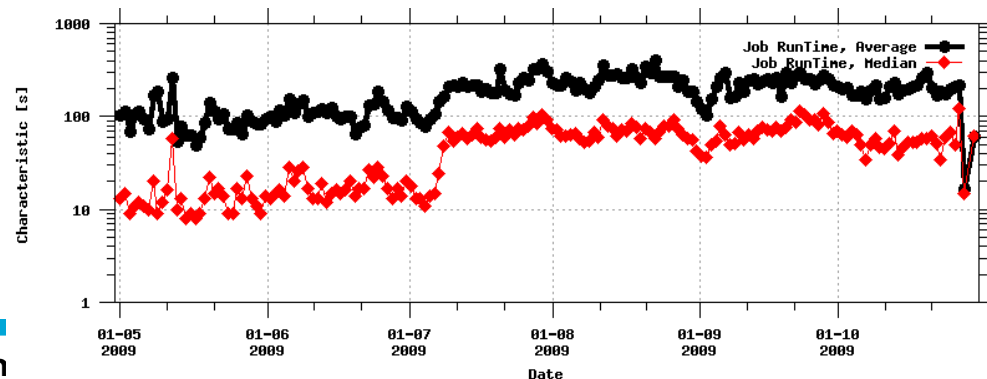
Trace ID	Total IO [MB]	Rd. [MB]	Wr [%]	HDFS Wr[MB]
CWA-01	10,934	6,805	38%	1,538
CWA-02	75,546	47,539	37%	8,563

- # Tasks/Job, ratio M:(M+R) Tasks, vary

- Understanding workload evolution



Noven



Agenda

1. Introduction to Cloud Computing
2. Is there a Role for Academics in Cloud Computing?
3. Which Applications Need Cloud Computing?
- 4. How to Build a Toolbox for Cloud Computing Research?**
5. How to Build the Core of Cloud Computing?
6. Conclusion

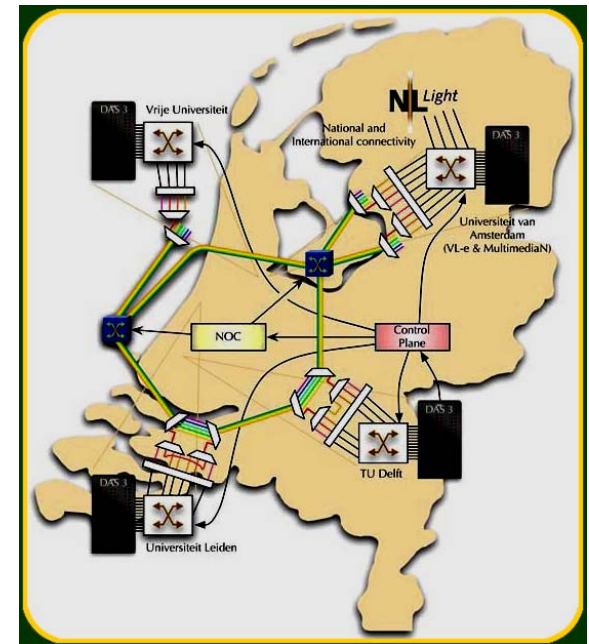
How to Build a Toolbox for Cloud Computing Research?

- Cloud Trace Archive
- Workload Models
- Large-scale experimental tools: simulators
- Benchmarking
- Service Failures

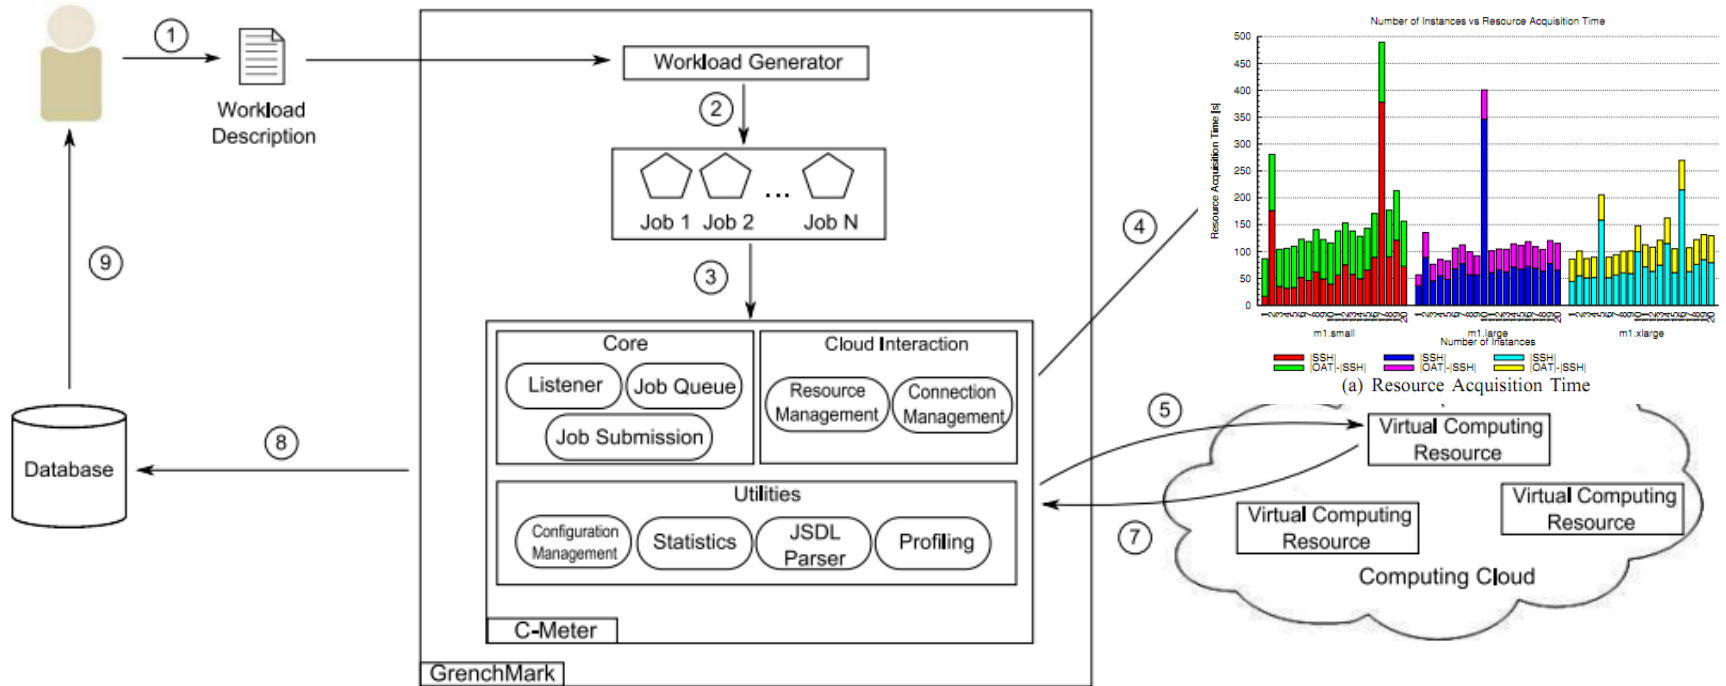
Research Opportunity: Large-Scale Experiments for Clouds

- The Infrastructure
 - Simulation
 - Emulation (DAS-4)
- The Data (Chicken and Egg Problem)
 - Operational logs of clouds
 - Cloud workload logs

ASCI
DAS-4



Research Opportunity: Cloud Benchmarking



- Which metrics? What is a good single-value indicator?
- IaaS: which provisioning and allocation policies?
- How to ensure reproducibility of the results?

Agenda

1. Introduction to Cloud Computing
2. Is there a Role for Academics in Cloud Computing?
3. Which Applications Need Cloud Computing?
4. How to Build a Toolbox for Cloud Computing Research?
- 5. How to Build the Core of Cloud Computing?**
6. Conclusion

At the Core: Deployment Model

- Using data centers for dynamic resource allocation




- Main advantages:
 1. Significantly lower over-provisioning
 2. Efficient coverage of the world is possible

At the Core: Performance Engineering



- **Many-Tasks Scientific Computing**
 - Identified proto-MT users from grid and PPEs
- **Performance Evaluation of Four Commercial Clouds**
 - Amazon EC2, GoGrid, Elastic Hosts, Mosso
 - Resource acquisition, Single- and Multi-Instance benchmarking
 - **Low compute and networking performance**
 - **Performance variability**
- **Clouds vs Other Environments**
 - Good for short-term, better perf. needed



1- Iosup et al., Performance Analysis of Cloud Computing Services for Many Tasks Scientific Computing, IEEE TPDS, 2011,

http://www.st.ewi.tudelft.nl/~iosup/cloud-perf10tpds_in-print.pdf

2- Iosup et al., On the Performance Variability of Production Cloud Services, CCGrid 2011, <pds.twi.tudelft.nl/reports/2010/PDS-2010-002.pdf>

Rich Research Opportunities: How to Build the Core?

- Deployment Models
- Scheduling
- Performance Engineering
- Reliability Engineering
- Scalability and Elasticity
- Cloud inter-operation (federation)
- Etc. (including Usability, Security, Utility Models, and Programming Models)

Agenda

1. Introduction to Cloud Computing
2. Is there a Role for Academics in Cloud Computing?
3. Which Applications Need Cloud Computing?
4. How to Build a Toolbox for Cloud Computing Research?
5. How to Build the Core of Cloud Computing?
- 6. Conclusion**

Research Opportunities in Cloud Computing

- **Understanding how real clouds work**
 - understand which applications work (in less than 10 years)
 - Massively social gaming = 250 million highly volatile customers
- **Building a toolbox for cloud computing research**
 - Modeling cloud infrastructure (performance, availability) and workloads
 - Compare clouds with other platforms (grids, parallel production env., p2p,...)
- **Building the core of cloud computing**
 - Deployment and Scheduling
 - Performance and reliability engineering
 - Scalability and elasticity
 - Cloud inter-operation
 - ... and many others



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

Thank you for your attention!

Questions? Suggestions? Observations?

More Info:

- <http://www.st.ewi.tudelft.nl/~iosup/research.html>
- http://www.st.ewi.tudelft.nl/~iosup/research_gaming.html
- http://www.st.ewi.tudelft.nl/~iosup/research_cloud.html

Alexandru Iosup

A.iosup@tudelft.nl

<http://www.pds.ewi.tudelft.nl/~iosup/> (or google "iosup")

Parallel and Distributed Systems Group
Delft University of Technology

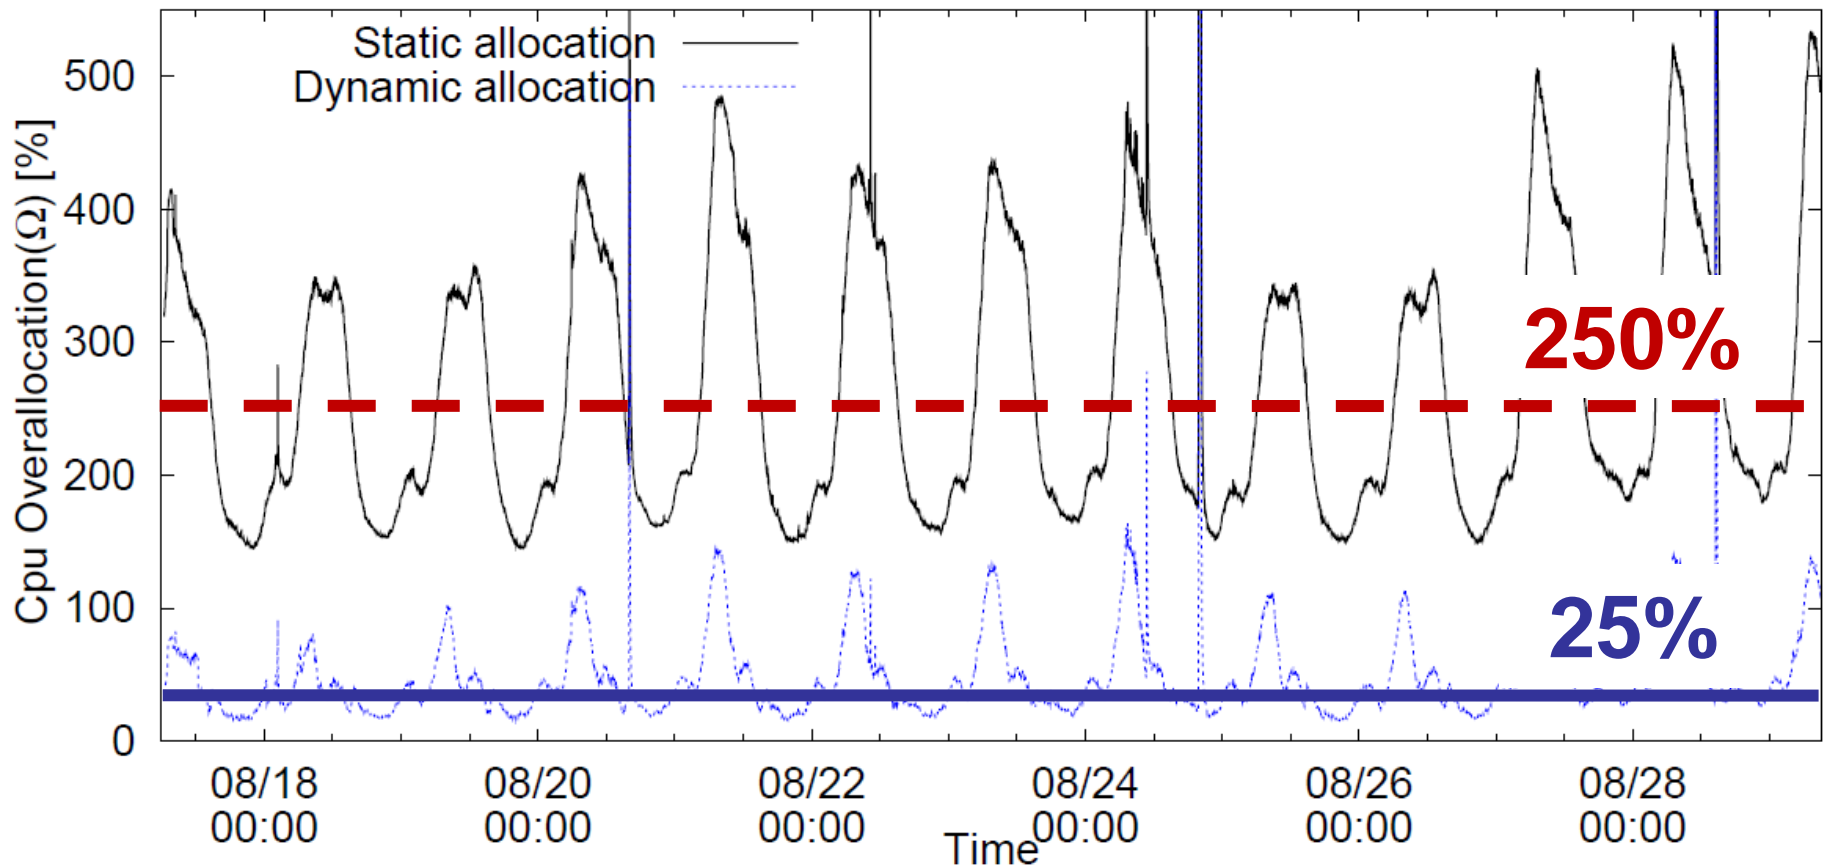
Do not hesitate to
contact me...



Additional Slides

Static vs. Dynamic Allocation


Q: What is the penalty for static vs. dynamic allocation?



Continuous Analytics for MMOGs

MMOG Data =


raw and derivative information
from the virtual world (millions
of users)



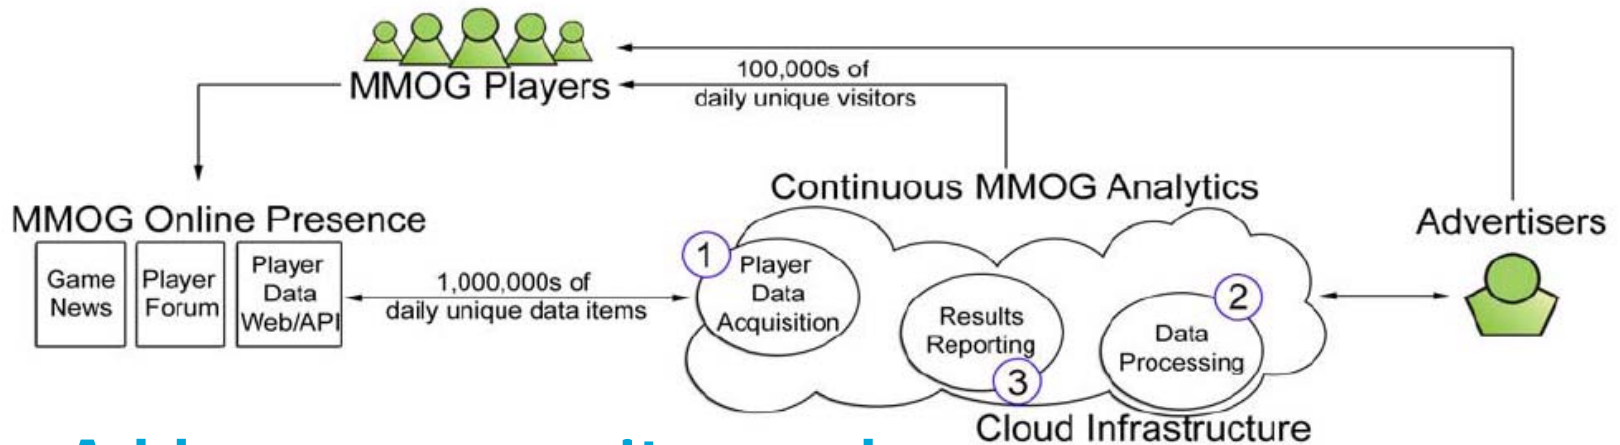
Continuous Analytics for MMOGs =

Analysis of MMOG data s.t.
important events are not lost

- Data collection
- Data storage
- Data analysis
- Data presentation
- ... at MMOG rate and scale



The CAMEO Framework*



1. Address community needs

- Can analyze skill level, experience points, rank
- Can assess community size dynamically

2. Using on-demand technology: Cloud Comp.

- Dynamic cloud resource allocation, Elastic IP

3. Data management and storage: Cloud Comp.

- Crawl + Store data in the cloud (best performance)


4. Performance, scalability, robustness: Cloud Comp.

* A. Iosup, CAMEO: Continuous Analytics for Massively Multiplayer Online Games on Cloud Resources. ROIA, Euro-Par 2009 Workshops, LNCS 6043, (2010)

CAMEO: Exploiting Cloud Features

- **Machines close(r) to server**

- Traffic dominated by small packets (latency)



- **Elastic IP to avoid traffic bans**


(legalese: acting on behalf of real people)

A. Iosup, A. Lascateu, N. Tapus, CAMEO: Enabling Social Networks for Massively Multiplayer Online Games through Continuous Analytics and Cloud Computing, ACM NetGames 2010.


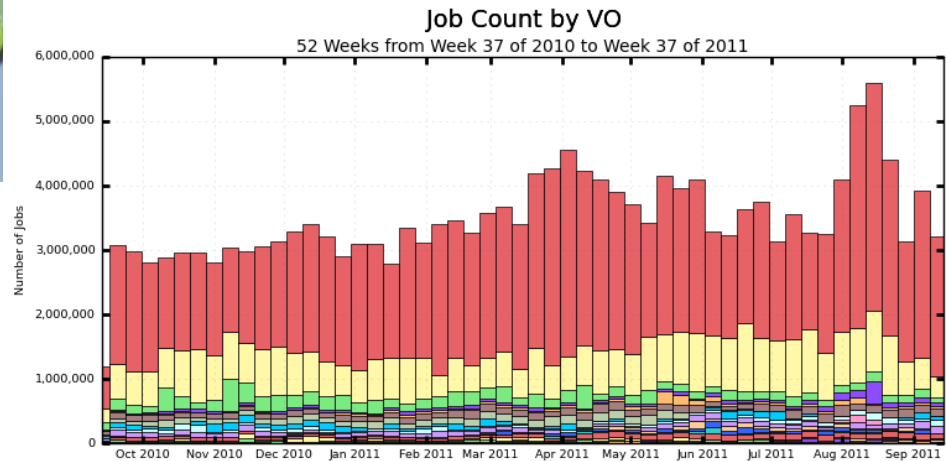
Sample Game Analytics Results

Skill Level Distribution in RuneScape

- **RuneScape**: 135M+ open accounts (largest)
- Dataset: **3M players (largest measurement, to date)**
 - 1,817,211 over level 100
 - Max skill 2,280
- **Put a price on MSG analytics**



... Avoid Rehashing Old Results



Maximum: 5,599,404 , Minimum: 1,198,436 , Average: 3,477,589 , Current: 3,208,759