



# LUND UNIVERSITY

## Cloud Control Workshop

Rantzer, Anders; Westin, Eva

2014

*Document Version:*

Publisher's PDF, also known as Version of record

[Link to publication](#)

*Citation for published version (APA):*

Rantzer, A., & Westin, E. (2014). *Cloud Control Workshop*. (Technical Reports TFRT-7639). Department of Automatic Control, Lund Institute of Technology, Lund University.

*Total number of authors:*

2

**General rights**

Unless other specific re-use rights are stated the following general rights apply:

Copyright and moral rights for the publications made accessible in the public portal are retained by the authors and/or other copyright owners and it is a condition of accessing publications that users recognise and abide by the legal requirements associated with these rights.

- Users may download and print one copy of any publication from the public portal for the purpose of private study or research.
- You may not further distribute the material or use it for any profit-making activity or commercial gain
- You may freely distribute the URL identifying the publication in the public portal

Read more about Creative commons licenses: <https://creativecommons.org/licenses/>

**Take down policy**

If you believe that this document breaches copyright please contact us providing details, and we will remove access to the work immediately and investigate your claim.

LUND UNIVERSITY

PO Box 117  
221 00 Lund  
+46 46-222 00 00



# Workshop: Cloud control

LUND CENTER FOR CONTROL OF COMPLEX ENGINEERING SYSTEMS  
LUND UNIVERSITY





# Cloud control workshop

7-9 May 2014

Old Bishop's Palace at Biskopsgatan 1 in Lund

## Scientific Committee

Maria Kihl, Lund University (chair)

Karl Erik Årzén, Lund University

Erik Elmroth, Umeå University

Tarek Abdelzaher, University of Illinois at Urbana-Champaign

Jie Liu, Microsoft Research

Bruno Sinopoli, Carnegie Mellon University

Vladimir Vlassov, Royal Institute of Technology

Giovanni Toffetti, IBM Haifa Research Lab

## Organizing Committee

Maria Kihl

Karl Erik Årzén

Erik Elmroth

Anders Rantzer

Eva Westin

**MAILING ADDRESS**

Department of Automatic Control  
Lund University  
Box 118  
SE-221 00 LUND, SWEDEN

**VISITING ADDRESS**

Institutionen för Reglerteknik  
Ole Römers väg 1  
232 63 LUND

**TELEPHONE**

+46 46 222 87 87

**FAX**

+46 46 13 81 18

**GENERIC E-MAIL ADDRESS**

control@control.lth.se

**WWW**

[www.lccc.lth.se](http://www.lccc.lth.se)

Printed: Media-Tryck, Lund, Sweden, Feb 2015

ISSN 0280-5316  
ISRN LUTFD2/TFRT--7639--SE

# Content

---

1. Introduction	5
1.1 Workshop Theme	5
1.2 Scope	5
1.3 Organization and venue	6
2. Panel discussion	6
2.1 Open problems (from an industry perspective) in Cloud control	6
2.2 Centralized vs. decentralized control for large-scale computing	6
2.3 Energy management for data centers	6
2.4 Resource optimization of bursty workloads in clouds	7
3. Panel discussion	8
Appendix A PROGRAM	9
Appendix B PARTICIPANTS	12
Appendix C PRESENTATIONS	14
Deploying complex applications to Google Cloud using Google Deployment Manager <b>Olia Kerzhner</b> , Google	14
Cloud Control Systems - Real Time Analytics <b>Simon Tuffts</b> , Econultant	19
Elasticity Manager for Elastic Key-Value Stores in the Cloud <b>Vladimir Vlassov</b> , Royal Institute of Technology	36
Principles and Methods for Elastic Computing <b>Schahram Dustdar</b> , Vienna University of Technology	47
Exploring Autonomics for Clouds <b>Manish Parashar</b> , Rutgers University	56
Thinking parallel: Multi-cores, virtual elasticity, and the application programmer <b>Geir Horn</b> , University of Oslo, Norway	65
Simplified Cloud Control Using Dimension Reduction <b>Jianguo Yao</b> , Shanghai Jiao Tong University	71
Real-time Performance Control of Elastic Virtualized Network Functions <b>Tommaso Cucinotta</b> , Bell Labs, Alcatel-Lucent, Ireland	79
An Adaptive Utilisation Accelerator for Virtualized Environments <b>Giovanni Toffetti</b> , IBM Haifa Research Lab	87
Dynamic Power Management in Data Centers: Theory & Practice <b>Mor Harchol-Balter</b> , Computer Science Department, Carnegie Mellon University	92
Reducing Power Delivery Costs for Cloud Data Centers <b>Bhuvan Urgaonkar</b> , Penn State University	100
Brownout: Building More Robust Cloud Applications <b>Cristian Klein</b> , Umeå University	115

RT-Xen: Real-Time Virtualization for the Cloud <b>Chenyang Lu</b> , Washington University in St. Louis	125
Service Level Agreement for Cloud Computing: Towards a Control-Theoretic Approach <b>Sara Bouchenak</b> , University of Grenoble	131
Performance-Energy Trade-off in Multi-Server Queueing Systems with Setup Delay <b>Samuli Aalto</b> , Aalto University	140
LCCC activities in Cloud Control <b>Maria Kihl</b> , Lund University	148
Capacity Management in IaaS Cloud <b>David Breitgand</b> , IBM Research Haifa	157
Deadline scheduling for big-data jobs and fault-tolerance of datacenter applications <b>Peret Bodik</b> , Microsoft Research	171
Control issues in warehouse-scale datacenters <b>John Wilkes</b> , Google	178
Application Performance Management in the Cloud using Learning, Optimization, and Control <b>Xiaoyun Zhu</b> , VMware Inc.	188
Event-based control: a way to reduce reconfiguration in autonomic computing <b>Nicholas Marchand</b> FGIPSA lab, France	208
Guided tour through a cloud datacenter – The Umeå University approach to cloud resource management <b>Erik Elmroth</b> , Umeå University	225

# 1. Introduction

---

LCCC workshops are organized in a 3-day format. About 20-25 speakers from academia and industry are invited for the workshop, selected for excellence and for an optimal coverage of the theme. The speakers are also encouraged to extend their stay beyond the workshop for further interaction with the local research environment. For each workshop, the research theme is chosen strategically to support the vision of a LCCC, usually with a cross-disciplinary perspective. An international scientific committee is responsible for the program.

## 1.1 WORKSHOP THEME

The Cloud Control Workshop was aimed to foster research in the multidisciplinary area of Cloud Control, leveraging expertise in areas such as distributed systems, control theory, autonomic computing, systems management, mathematical statistics, energy management, performance management, etc, to manage the cloud. The aim was to gather researchers from both academia and industry, and thereby promote new collaborations and ideas.

## 1.2 SCOPE

The workshop addressed challenges regarding the management and control of large-scale cloud infrastructures by bringing together the computer science community doing cloud management with the control community dealing with large-scale computing systems.

## 1.3 ORGANIZATION AND VENUE

The workshop was initiated by Maria Kihl (Dept. of Electrical and Information Technology, Lund University), Karl Erik Årzén (Dept. of Automatic Control, Lund University) and Erik Elmroth (Dept. of Computer Science, Umeå University).

The scientific committee consisted of Maria Kihl (chair), Karl Erik Årzén, Erik Elmroth, Tarek

Abdelzaher (University of Illinois at Urbana-Champaign), Jie Liu (Microsoft Research), Bruno Sinopoli (Carnegie Mellon University), Vladimir Vlassov (Royal Institute of Technology), and Giovanni Toffetti (IBM Haifa Research Lab).

The local organization and interactions with workshop speakers and participants was handled by Eva Westin.

The workshop was held at the Bishop's Palace at Lund University, May 7-9 2014.

## 2. Group discussions

---

During the workshop, there were four group discussions on important issues in Cloud control. A short summary of each discussion follows.

### **2.1 OPEN PROBLEMS (FROM AN INDUSTRY PERSPECTIVE) IN CLOUD CONTROL**

In this discussion, the focus was on research problems in Cloud control that are particularly interesting from an industry perspective.

The discussion was mainly focused on energy, which was believed as an important issue for industry. However, several of the industry representatives did not believe that turning off physical machines is a feasible solution for energy optimization. Instead, moving jobs to servers with low loads may be a better solution. Some types of jobs, for example big data jobs, can be placed wherever there is spare capacity. Also, real-time power optimization introduces one more complexity that may not be optimal for the system. Further, VMWare for example, has a solution for power reduction, but few customers use it. Also, there has been some work on virtualized batteries, which would be a way to make energy consumption sellable, that is, customers may buy a certain amount of power for their jobs.

Further, there was some discussion on application aware placement and infinite cloud architectures. For example, in Sweden, telecom operators are moving base station intelligence in to the cloud. Another example is cloud gaming applications. This will require new architectures with application-aware placement of components, which is a difficult challenge to solve.

### **2.2 CENTRALIZED VS. DECENTRALIZED CONTROL FOR LARGE-SCALE COMPUTING**

This discussion was focused on issues related to centralized vs. decentralized control for large-scale computing. One general conclusion was that there is no single rule to apply on all of the problems. Also, when you want to design a system to either be centralized or decentralized, different aspects and trade-offs need to be considered such as flexibility vs. agility and robustness, and scalability vs. optimality.

Generally speaking, the problems that we are dealing with in a cloud resource management domain are in principal so complex that centralized solutions will have their limitations no matter what. Instead, a feasible point of optimality should be found, and find decentralized solutions that benefits, for example, scalability and flexibility.

It was also mentioned that at some level of abstraction every system is naturally decentralized. The suggested solution should be a little bit of centralization and a lot of decentralization. From a design perspective, computation is better to be as decentralized as possible, while the data can be stored centrally.

### **2.3 ENERGY MANAGEMENT FOR DATA CENTERS**

This discussion focused on several issues related to energy management for data centers. Much of the discussion was once more focused on the question of turning off and on machines. Also, it was discussed how exhaust energy and heat from data centers can be stored using batteries and water, respectively. The saved energy can be used to power surges in demand, and to flatten the energy usage curve. Additionally, there are

examples of where exhaust energy has been put to good use, heating adjacent homes, facilities, and waterbeds. However, taxation can make it unprofitable to sell exhaust energy.

It was agreed that data centers are not an efficient mean to generate heat and should not be considered a part of the energy infrastructure. Rather, effort should be directed towards reducing energy consumption through heat profiling and heat-aware placement. The discussion of VM placement and migration as a heat optimization parameter, and how challenging scaling down is, was reignited. Although it was agreed that migrating VMs would most likely require more energy in the proportion to the heat gained.

From a data center management perspective, the discussion came to evolve around the fact that different applications have different energy profiles that do not necessarily translate into what they pay for the service, depending on the type of workload, time of day, and the type of energy used. Thereafter the discussion was focused on how energy costs are transferred to the customer, and how application administrators are primarily concerned with overall performance. From where the session ended with a discussion on various energy saving schemes, ranging from energy aware hardware in the mobile industry to energy pricing profiles. It was proposed that energy should be provisioned and purchased separately by the application developer, but it was agreed that such a model would inhibit data-center flexibility and distract from the cloud paradigm.

## 2.4 RESOURCE OPTIMIZATION OF BURSTY WORKLOADS IN CLOUDS

This discussion was focused on topics related to resource optimization in clouds when there are bursty workloads with spikes. First of all, it is not easy to define a bursty workload. Load can be measured in several ways and requests can generate very different amount of processing. Also, bursts come at different time-scales and,

therefore, it is not feasible to ask an "oracle" every millisecond. Interesting bursts are the ones that cannot be handled by "ordinary" control. Proactive (predictive) techniques should be used to detect spikes. Queuing theory could be used to analyze bursts, however few people use it and the question is how accurate M/M/Q models are. One input is that it is better to do something based on M/M/Q queues than just use ad-hoc methods.

Also, the topic of centralized vs. decentralized control came back. Centralized control should be used as much as possible, and then be complemented with decentralized control.

### 3. Panel Discussion on "Challenges for Cloud control

---

The workshop ended with a panel discussion on "Challenges for Cloud control". The members of the panel were John Wilkes, Simon Tuffs, David Breitgand, and Geir Horn. A short summary of the discussion topics is give here.

Most of the discussion was focused on what challenges and problems the universities, and in particular PhD students should work on. A general conclusion was to always reach out to industry and solve real problems. One argument was that universities should work on fundamental problems that industry does not have time to solve. Availability is really important. Optimal solutions are not so important, better to have suboptimal solutions that are scalable and improve availability.

Another challenge is the lack of real data. Companies do not share data. One solution for this is to buy cloud capacity (which is very cheap) and run experiments directly on the cloud. However, this will not work for all research topics, for example scheduling. Another solution is to build your own cloud and bombard it with traffic, real or synthetic.

Part of the discussion was devoted to the notion of "Cloud". One argument was that the term "Cloud" will disappear in a few years, instead everything will be called "Big Data" or something else. However, the fundamental challenges will remain so the future is bright.

# Appendix A

## PROGRAM

---

Wednesday, 7 May

- 09:30 Registration and coffee
- 10:00 Opening — Maria Kihl
- 10:15 Deploying complex applications to Google Cloud using Google Deployment Manager  
Olia Kerzhner, Google  
Cloud Control Systems - Real Time Analytics  
Simon Tuffs, consultant  
Energy and Resource Management Challenges in Data Centers  
Tarek Abdelzaher, University of Illinois at Urbana Champaign
- 12:00 Lunch
- 13:00 Group Discussions
- 13:30 Group Presentation
- 14:00 Elasticity Manager for Elastic Key-Value Stores in the Cloud  
Vladimir Vlassov, Royal Institute of Technology  
Principles and Methods for Elastic Computing  
Schahram Dustdar, Vienna University of Technology
- 15:00 Coffee
- 15:30 Exploring Autonomics for Clouds  
Manish Parashar, Rutgers University  
Thinking parallel: Multi-cores, virtual elasticity, and the application programmer  
Geir Horn, University of Oslo, Norway  
Simplified Cloud Control Using Dimension Reduction  
Jianguo Yao, Shanghai Jiao Tong University



Thursday, 8 May

- 09:00 Real-time Performance Control of Elastic Virtualized Network Functions  
Tommaso Cucinotta, Bell Labs, Alcatel-Lucent, Ireland  
An Adaptive Utilisation Accelerator for Virtualized Environments  
Giovanni Toffetti, IBM Haifa Research Lab
- 10:00 Coffee
- 10:30 Dynamic Power Management in Data Centers: Theory & Practice  
Mor Harchol-Balter, Computer Science Department, Carnegie Mellon University  
Reducing Power Delivery Costs for Cloud Data Centers  
Bhuvan Urgaonkar, Penn State University  
Brownout: Building More Robust Cloud Applications  
Cristian Klein, Umeå University
- 12:00 Lunch
- 13:00 Group Discussions
- 13:30 Group Presentations
- 14:00 RT-Xen: Real-Time Virtualization for the Cloud  
Chenyang Lu, Washington University in St. Louis  
Service Level Agreement for Cloud Computing: Towards a Control-Theoretic Approach  
Sara Bouchenak, University of Grenoble
- 15:00 Coffee
- 15:30 Performance-Energy Trade-off in Multi-Server Queueing Systems with Setup Delay  
Samuli Aalto, Aalto University
- 17:30 Buses leave from Bangatan (green arrow on map)
- 18:30 Workshop Dinner at Turning Torso, Malmö



Friday, 9 May

- 09:00 LCCC activities in Cloud Control  
Maria Kihl, Lund University  
Capacity Management in IaaS Cloud  
David Breitgand, IBM Research Haifa
- 10:00 Coffee
- 10:30 Deadline scheduling for big-data jobs and fault-tolerance of datacenter applications  
Peter Bodik, Microsoft Research  
Control issues in warehouse-scale datacenters  
John Wilkes, Google  
Application Performance Management in the Cloud using Learning, Optimization, and Control  
Xiaoyun Zhu, VMware Inc.
- 12:00 Lunch
- 13:00 Modern Infrastructure: The Convergence of Network, Compute, and Data  
Jason Hoffman, Ericsson  
Event-based control: a way to reduce reconfiguration in autonomic computing  
Nicholas Marchand, GIPSA lab, France  
Guided tour through a cloud datacenter - The Umeå University approach to cloud resource management  
Erik Elmroth, Umeå University
- 14:30 Coffee
- 15:00 Panel Discussion: Challenges in Cloud Control  
15:30 Final Remarks and End of Workshop



# Appendix B

## PARTICIPANTS LCCC Focus Period

### Workshop on Cloud Control, May 2014

---

Aalto, Samuli	Aalto University	samuli.aalto@aalto.fi
Abdelzaher, Tarek	University of Illinois at U-C	zaher@illinois.edu
Ali Eldin Hassan, Ahmed	Umeå University	ahmeda@cs.umu.se
Årzén, Karl-Erik	Lund University	karlerik@control.lth.se
Åström Karl-Johan	Lund University	kja@control.lth.se
Berekmeri, Mihaly	GIPSA-lab, Grenoble	mihaly.berekmeri@gipsa-lab.grenoble-inp.fr
Bini, Enrico	Scuola Superiore St'Ana, Pisa	e.bini@sssup.it
Blomdell, Anders	Lund University	anders.blomdell@control.lth.se
Bodik, Peter	Microsoft	peterb@microsoft.com
Bouchenak, Sara	IMAG	sara.bouchenak@imag.fr
Breitgand, David	IBM	davidbr@il.ibm.com
Como, Giacomo	Lund University	giacomo.como@control.lth.se
Cucinotta Tommaso	Bell Labs, Alcatel-Lucent	tommaso.cucinotta@alcatel-lucent.com
Dellkrantz, Manfred	Lund University	manfred.dellkrantz@control.lth.se
Dustdar, Schahram	TU Wien	dustdar@infosys.tuwien.ac.at
Dürango, Jonas	Lund University	jonas.durango@control.lth.se
Eker, Johan	Lund University / Ericsson	johan.eker@control.lth.se
Elmroth, Erik	Umeå University	elmroth@cs.umu.se
Gambi, Alessio	Università della Svizzera Italiana	alessio.gambi@usi.ch
Gran, Ernst Gunnar	Simula Research Laboratory	ernstgr@simula.no
Harchol-Balter, Mor	Carnegie Mellon University	harchol@cs.cmu.edu
Hernandez, Francisco	Umeå University	francisco@cs.umu.se
Hoffman, Jason	Ericsson	jason.a.hoffman@ericsson.com
Horn, Geir	University of Oslo	geir.horn@mn.uio.no
Ibidunmoye, Olumuyiwa	Umeå University	muyi@cs.umu.se
Kerzhner, Olia	Google	olia@google.com
Kihl, Maria	Lund University	maria.kihl@eit.lth.se
Klein, Cristian	Umeå University	cristian.klein@cs.umu.se
Krywda, Jakub	Umeå University	jakub@cs.umu.se
Landfeldt, Björn	Lund University	bjorn.landfeldt@eit.lth.se
Ljung, Peter	Sony Mobile	peter.ljung@sonymobile.com
Lorido-Botran, Tania	University of the Basque Country/ Umeå Univ.	tbotran@cs.umu.se
Lu, Chenyang	Washington University in St Louis	lu@cse.wustl.edu
Madjidian, Daria	Lund University	daria.madjidian@control.lth.se
Maggio, Martina	Lund University	martina.maggio@control.lth.se
Marchand, Nicolas	GIPSA-lab	nicolas.marchand@gipsa-lab.fr
Mehta, Amardeep	Umeå University	amardeep@cs.umu.se
Nilsson, Anders	Lund University	anders.nilsson@control.lth.se
Östberg, P-O	Umeå University	p-o@cs.umu.se

Papadopoulos, Alessandro	Lund University	alessandro.papadopoulos@control.lth.se
Parashar, Manish	Rutgers University	parashar@rutgers.edu
Rantzer, Anders	Lund University	anders.rantzer@control.lth.se
Robertsson, Anders	Lund University	anders.robertsson@control.lth.se
Robu, Bogdan	GIPSA-lab	bogdan.robu@gipsa-lab.grenoble-inp.fr
Sedaghat, Mina	Umeå University	mina@cs.umu.se
Skeie, Tor	Simula Research Laboratory/ Univ. of Oslo	tskeie@simula.no
Tärneberg, William	Lund University	william.tarneberg@eit.lth.se
Toffetti, Giovanni	IBM Haifa Research Lab	giovanni@il.ibm.com
Tomas, Luis	Umeå University	luis@cs.umu.se
Tordsson, Johan	Umeå University	tordsson@cs.umu.se
Truong, Hong-Linh	TU Wien	truong@dsg.tuwien.ac.at
Tuffs, Simon	consultant	simontuffs@gmail.com
Urgaonkar, Bhuvan	Penn State University	bhuvan@cse.psu.edu
Vlassov, Vladimir	KTH, Stockholm	vladv@kth.se
Wang, Cheng	Penn State University	cxw967@cse.psu.edu
Westin, Eva	Lund University	eva.westin@control.lth.se
Wittenmark, Björn	Lund University	bjorn.wittenmark@control.lth.se
Wilkes, John	Google	johnwilkes@google.com
Yao, Jianguo	Shanghai Jiao Tong University	jianguo.yao@sjtu.edu.cn
Zhu, Xiaoyun	VMware	xzhu@vmware.com



# Appendix C

## PRESENTATIONS

---

### DEPLOYING COMPLEX APPLICATIONS TO GOOGLE CLOUD USING GOOGLE DEPLOYMENT MANAGER

Olia Kerzhner, Google

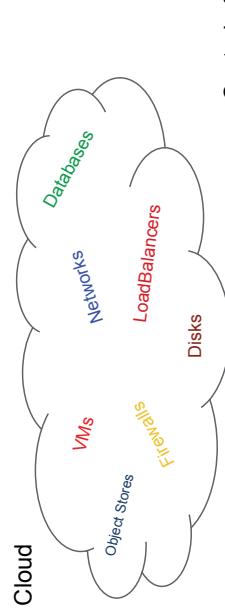
Deploying dynamic distributed applications in the cloud is hard. Google Deployment Manager is the new Google Cloud service that provides a simple templating mechanism that allows you to declaratively describe your solution, and then deploy it with a single command. Deployment Manager then provisions, scales, and monitors your solution. In this talk I will introduce Google Deployment Manager and demonstrate how it can be used to easily and repeatably declare and deploy dynamic systems in Google Cloud





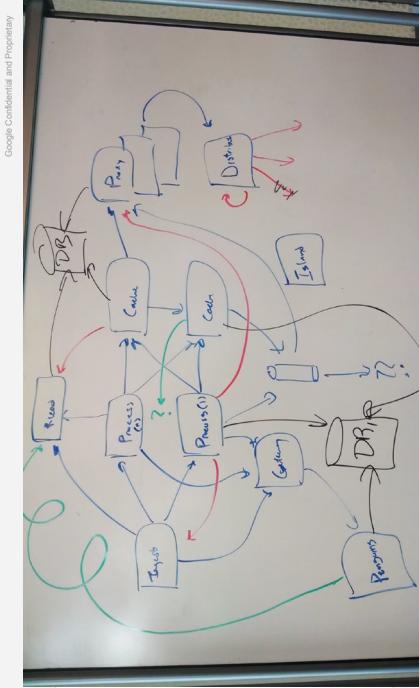
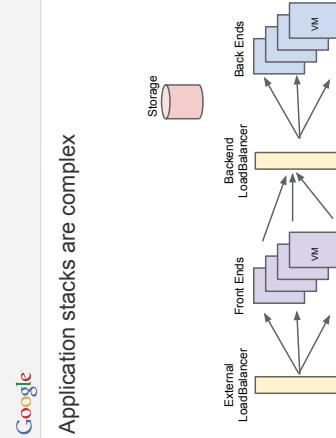
## Deploying complex applications to Google Cloud

Olia Kerzhner  
olia@google.com

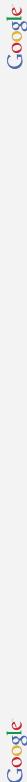


Control ..?

Application stacks are complex



Google Confidential and Proprietary

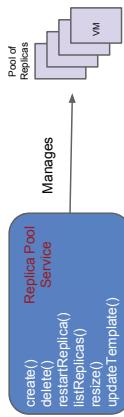


## Building on top of Google Compute Engine APIs

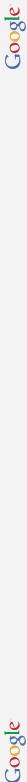
- GCE API are great for manipulating individual resources
- any minimally complex application requires users to write scripts or other code to deploy
- using scripts or manually creating resources results in "snowflake" deployments
- many open-source or proprietary tools try to solve this problem

## Introducing: Replica Pool API

- Powerful new API to manage sets of homogenous VMs
  - VMs are created based on a template
  - declaratively specify what software to deploy
  - VMs are monitored for health and are restarted on failed healthchecks
  - the API supports the `resize()` operation that will add or subtract VMs to/from the set



Google Confidential and Proprietary

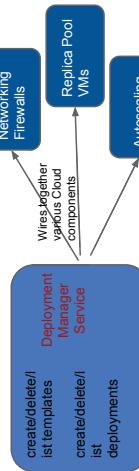


## Managing more complexity

- need multiple sets of homogenous VMs: FrontEnds, BackEnds, Batch workers
- need load balancers, firewalls, networks, etc.
- need a way to tie the components together

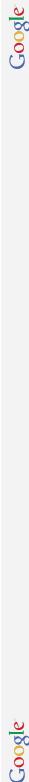
## Introducing: Deployment Manager

- a single declarative JSON template to tie together different Cloud components to define an application stack
- templates are reusable, overrides are supported for customization of deployments



Google Confidential and Proprietary

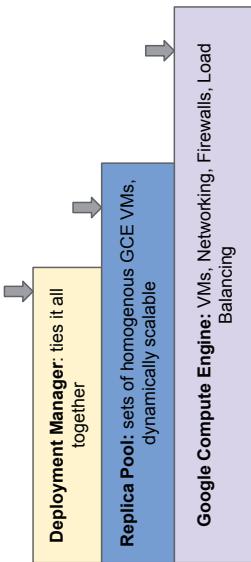
Google Confidential and Proprietary



## Deployment Template

- A DECLARATIVE specification for how to deploy a set of Google Cloud resources together into an application stack
- defined in terms of "modules", each module corresponds to a Google Cloud API
- modules are connected together via module names and references

## Google Cloud Layered APIs



Google Confidential and Proprietary



## Building Cloud-aware applications

What if you built your clustered application with the knowledge that it will run in the Cloud?

Traditional Cluster setup approaches:

- a hard-coded list of nodes supplied in the cluster on initialization
- connect peer-to-peer for discovery (complex configuration for peer connections)
- designating a master for set management
- brittle, complex, not scalable

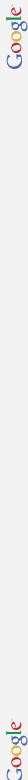
Google Confidential and Proprietary



## Enabling Cloud-aware cluster setup

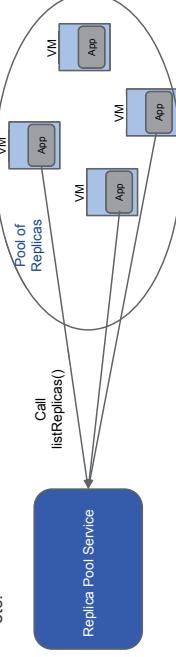
Google Confidential and Proprietary

Google Confidential and Proprietary



## Replica Pool provides Set membership discovery

- there is now a REST endpoint that lists members (replicas) in a set
- adjusted dynamically as the set grows or shrinks
- Cloud-aware applications can use it to set up clusters, find peers, shard, etc.



Give it a try!

- released to Limited Preview at the Google Cloud event in March
- documentation is public, usage requires whitelisting
- <https://developers.google.com/deployment-manager>
- <https://developers.google.com/compute/docs/replica-pool>

Try the sample templates, try writing your own. Send us your feedback!

## Summary

- **Replica Pool** is used to deploy scalable sets of homogenous VMs
- **Deployment Manager** is used to deploy and tie together multiple Replica Pools and other Cloud resources
- Setting up a cluster using Replica Pool allows for easy and robust cluster setup

Google Confidential and Proprietary



Google Confidential and Proprietary

Google Confidential and Proprietary

**CLOUD CONTROL SYSTEMS - REAL TIME ANALYTICS****Simon Tuffs, consultant**

Event Driven, Service-Oriented Cloud Systems can be viewed as systems of multi-variable time-series data. This makes them amenable to analysis using discrete-time signal-processing and feedback control systems analytic techniques. Some illustrative examples from a real-world Cloud System are presented: Anomaly detection and classification using statistical tests, correlation, and causal inference based on service dependencies; Qualification of new code deployments into production, automated failure detection and recovery; and Auto-scaling challenges and remedies, with use of feed-forward predictors to survive outages. A general architecture for Cloud Systems Analytics is presented, together with a view of the interesting challenges ahead.



# Cloud Control Systems - Real-Time Analytics

Dr. Simon Tuffs  
Cloudstream

<http://tinyurl.com/qafuyzn2>

## This Talk

- Cloud Based Service Applications
  - Analytics & Control, how and why
- Cloud Applications As:
  - Multi-variable time-series systems
  - Amenable to signal processing
  - Resilient to failure using analytics
  - Operational using feedback control

## Background

- University of Oxford
  - Self Tuning Control Systems/GPC
- 25 years software industry: highlights
  - Iridium Ground Station (Motorola)
  - Spacestation Infrastructure (Boeing)
  - Cloud (Netflix)
- Theme: Software at Scale

## Cloud Application Architectures

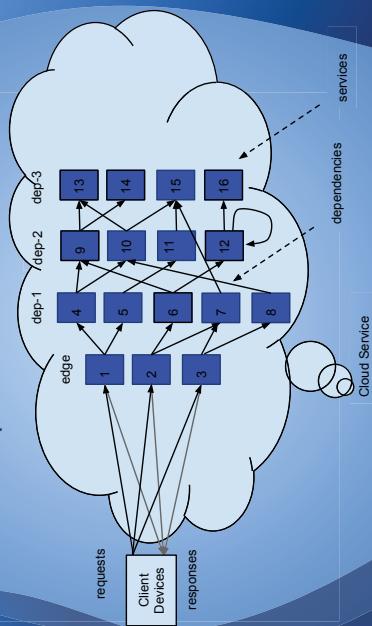
## Cloud: Application Architectures

- Classes
- Micro-service
- Massive volume business operations (e.g. Netflix)
- Big-Data
  - Terabytes of data, captured from streams into persistent stores.
  - Sparse compute intensive operations

## Cloud: Application Architectures

- Cloud Application composed of services
- Services have dependencies (graph)
- Requests flow from the edge down
- Responses flow back to edge
- Latencies accumulate forwards
- Errors propagate backwards
- Services are developed independently

## Cloud: Application Architectures Services & Dependencies:

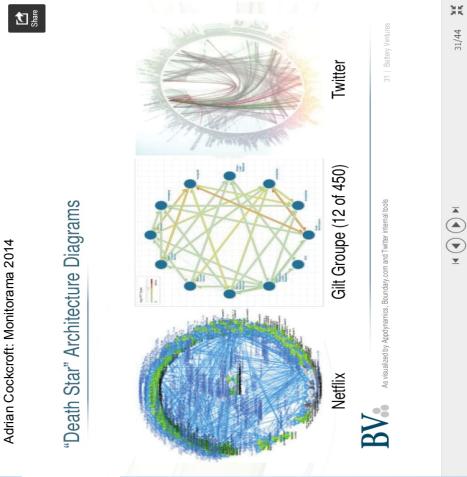


## Cloud: Application Architectures



**"We are not alone"**

"Death Star" Architecture Diagrams



## Cloud: Application Architectures

- Complex beyond human comprehension
- Nonlinear
- Time-varying
- Partially predictable
- Potentially chaotic
- The worst kind of "system"

## Cloud Applications: Analytics Classes

- Operational
  - Availability, fault detection, repair, performance optimization
- Business Intelligence
  - how much money are we making?
  - how many customers did we just lose?
  - how can we make more money?

**Analytics are not optional,  
they are essential**

## Cloud Applications: Monitoring

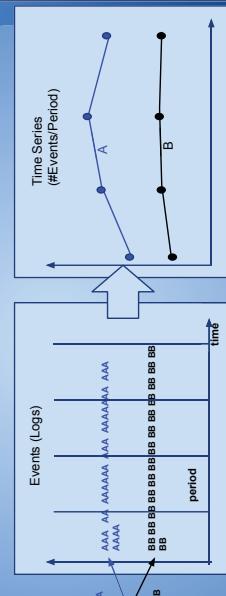
- To analyze you must monitor
- How do you handle billions of events?
- How do you transform them for analytics?

## Cloud Applications: Monitoring

- Instrument services:
  - to expose internal details (e.g. type of errors, versus HTTP 503's)
- With significant request volume:
  - monitored events become statistically driven time-series
  - signal processing methods then apply

## Cloud: Monitoring

- From Events To Time Series:

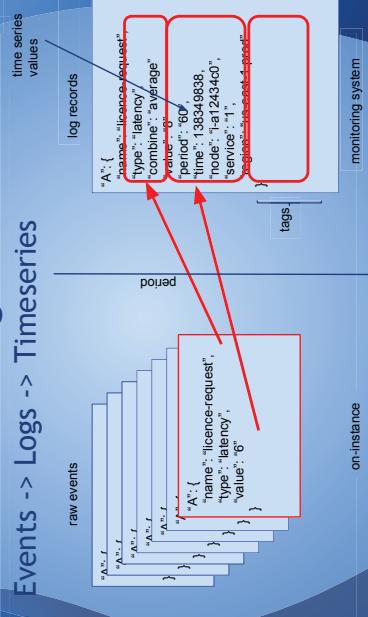


## Cloud: Monitoring Architecture

- Convert events to time-series (coordinate transform)
  - bucket by period
  - classify & tag
  - store for query/retrieval
- Reduces dimension of data by many orders of magnitude
  - -> Real Time Analytics become feasible

## Cloud: Monitoring Architecture

- Events -> Logs -> Timeseries



## What to Monitor?

- “Assume that any metrics not being analyzed will turn out to be garbage”
  - Adrian Cockcroft, Architect Netflix Cloud
- Instrument to measure:
  - health (success, failure)
  - performance (load, cpu)
  - availability (timeouts, fallbacks)
  - resources (disk i/o, memory, handles),
  - sla's (latency)

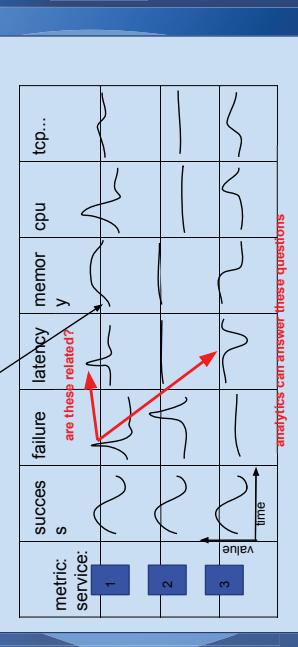
## Service Metric Visualization

- Classify metrics by type
- View services as rows of service:metrics
- Patterns start to emerge between visually.
- This scales to 100's of services and metrics (make the graphs small, human visual cortex sees patterns)

## Visualization as an Analytic

## Cloud: Visualizing service:metric

- 



## Beyond Visualization: Computational Analytics

analytics can answer these questions

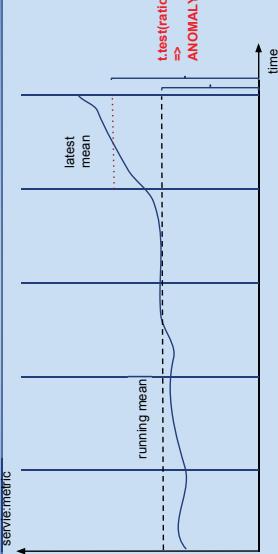
## Anomaly Detection

- Look at a service:metric
- Is it behaving normally, or is it showing signs of distress?
- How can we automate this?
- Without lots of configuration?
- In a scale invariant way?
- Use a mean-shift analytic...

## Anomaly Detection & Diagnosis

## Analytics for Anomalies?

- mean? variance?



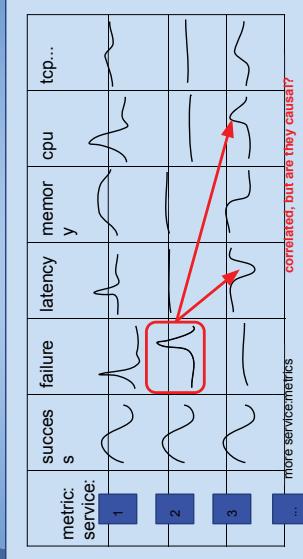
## Analytics for Anomalies?

- You found an anomalous service:metric, now what?

- Correlate against \*all other\* service metrics
- This is fast (<0.1s for 400sm in R)

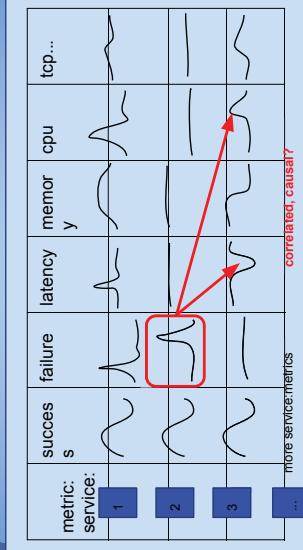
## Correlate

- Pearson + mean removal



## Filter

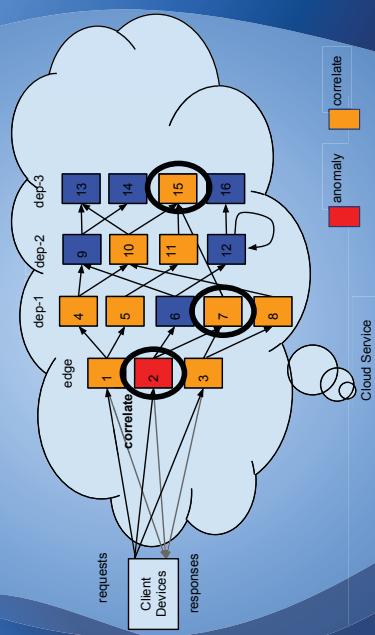
- Increase signal-to-noise:



## Can we do more?

- Correlation x Dependency = Probable Cause

## Anomaly -> Correlation -> Cause

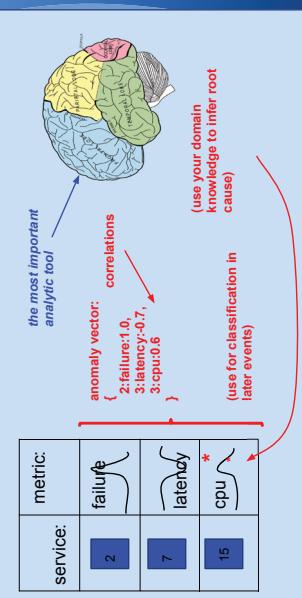


## Classify and Decide.

- Prune with dependency tree

## Build a model

- Persist this pattern for future causal analysis
- Did we see this anomaly vector before?

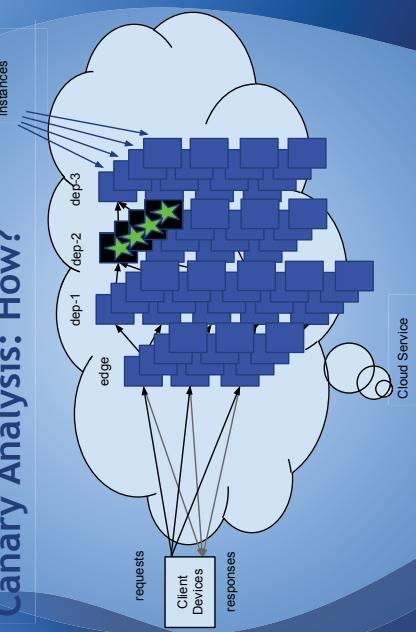


## Canary Analysis Defined

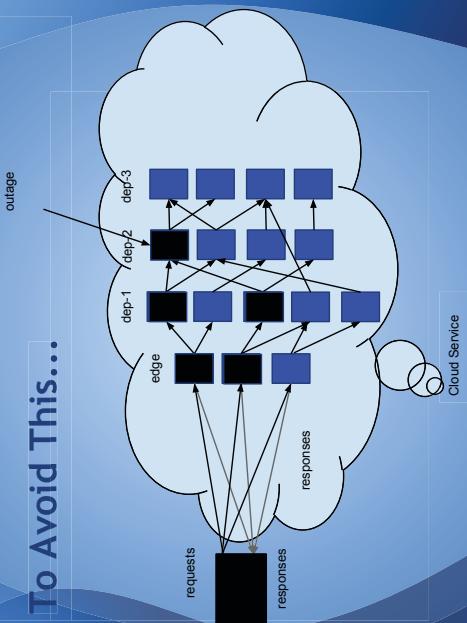
- For a given service:
  - Deploy new code to limited #instances
  - Analyze against existing production code
  - Decide whether good or bad
  - Push forward (upgrade all service instances)
    - or roll back.

## Canary Analysis (deployment)

## Canary Analysis: How?



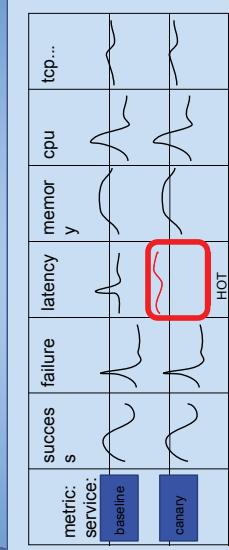
## To Avoid This...



## Canary Analysis

- How does this work?
  - Service metric grid (again), 2 rows.
  - Compare canary to baseline, statistical tests.

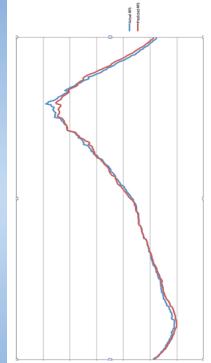
## Automated Canary Analysis



## Autoscaling

### Load Based Autoscaling

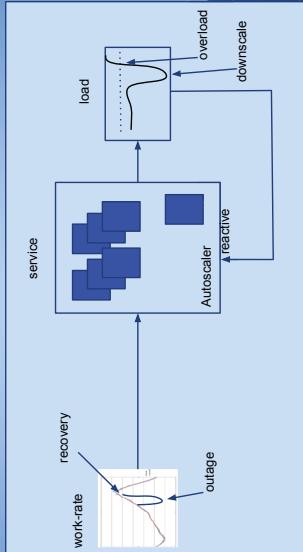
- increase #instances when load increases
- decrease #instances when load decreases
- works well...



## Except when it doesn't..

- During an outage, load drops
- Instances are terminated
- Service becomes underprovisioned for return to normal request rate
- Overload occurs
- Other services suffer.
- Chaos.

## Reactive Autoscaling

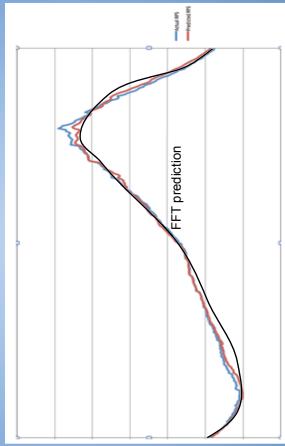


## How do you avoid this?

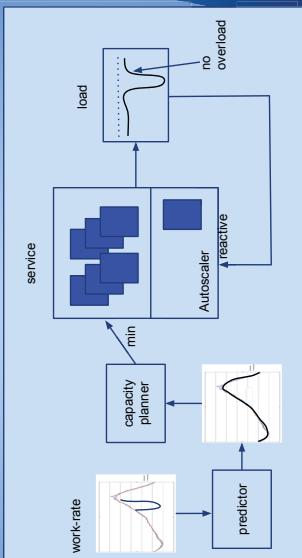
- Use feedforward control
- Base on prediction of request rate
- Simple application of FFT low-pass filter.

## Scryer

- FFT based prediction



- Netflix: Scryer
- Predictive+Reactive = Feedback Control



## Real Time Analytics Engine

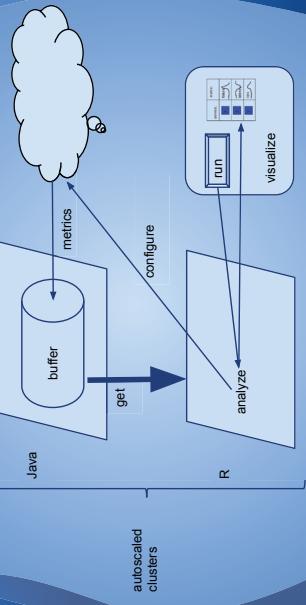
- One possible architecture: Java and R engines in the Cloud
  - gathering data
  - running analytics
  - performing visualization
  - doing notification

## Analytics at Scale

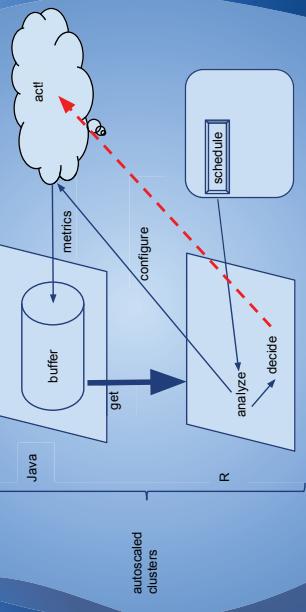
- How do you do analytics at scale?
  - Do monitoring at scale
  - Do data-collection & buffering at scale
  - Run Analytics at scale
  - Use the Cloud to achieve scale.  
(But use a different Cloud).

## Analytics at Scale

## Cloud Analytics: Interactive



## Cloud Analytics: Automated



## Cloud Analytics: Big Challenges

- instance outlier detection at scale
- tuning queues & timeouts for services
- detection of overload/underprovision
- anomaly detection (prediction)
- behavior pattern classification
- automatic alert tuning
- “closing the loop”

## Analytics Challenges



## Cloudstream

<https://github.com/simontuffs/cloudstream/wiki>

- Cloudstream Stack:

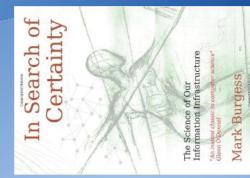
- Netflix OSS, Open/CPU, iPython, Cloudsim, Amazon/Kinetics Netflix/Suro Storm/Spark
- Real-Time Analytics, in the Cloud, for the Cloud.
  - Currently building an application simulator
  - Design & train analytics

## “Cloudstream”

## Questions?



## Recommendation:



- Mark Burgess
  - In Search Of Certainty, 2013
  - Views information systems from a physics perspective, showing the non-deterministic complexity we are creating, and how hard it is to manage

## Caution!

- Please seek a second opinion before spending years building a Ph.D. out of the following speculations & observations....

## A Posteriori Observations

- Focus on \$ not KWh for allocation
  - (they are isomorphic)
  - \$ drive customer behavior the right direction
- Consider standardizing on “Model Predictive Controls” (e.g. GPC)
  - Superset all other linear methods, save time :)
- Most of my challenges do not close any control loops
  - other than estimation/modeling loops

## A Posteriori Challenges

- Monitoring Validation
  - Our Cloud is down! Our Monitoring is down!
  - How can you tell?
- Avoid WOM (write-only monitoring)
  - how to aggregate useful data without losing information but still do analytics
- Causality
  - Infer dependency graph from data?
  - Cross-covariance for causation.

## A Posteriori Challenges

- Develop Cloud invariants / assertions as “models of behavior”
  - increased latency => upstream errors
  - upstream errors => downstream request drop
  - increased cpu => increased latency
  - increased requests => increased (cpu, load)
  - parameterize & tune a behavioral model base on these invariants.

## A Posteriori Challenges

- Machine learning (SVM, markov models)
  - Behavioral classification
  - Failure identification
- Evidence based learning
  - Bayesian networks for fault detection.
- Better predictors
  - Wavelets, basis functions.
- Modeling the Cloud
  - Dynamic Equilibrium
  - Transient Dynamics
  - “Kalman” Filtering

## A Posteriori Challenges

- Auto-tune configuration parameters  
*(close the loop)*
  - 99.5% latency  $\Leftrightarrow$  errors => need to increase caller timeouts.
  - 99.5% latency  $\Leftrightarrow$  load => need to scale up if at the “knee”.
  - 99.5% queue size ~ max-size => need to add worker threads
  - do this in production, across operating ranges

Thankyou!

**ELASTICITY MANAGER FOR ELASTIC KEY-VALUE STORES  
IN THE CLOUD****Vladimir Vlassov, Royal Institute of Technology**

The increasing spread of elastic Cloud services, together with the pay-as-you-go pricing model of Cloud computing, has led to the need of an elasticity controller. The controller automatically resizes an elastic service in response to changes in workload, in order to meet Service Level Objectives (SLOs) at a reduced cost. However, variable performance of Cloud Virtual Machines and nonlinearities in Cloud services, such as the diminishing reward of adding a service instance with increasing the scale, complicates the controller design. First, we briefly discuss challenges and some approaches to automation of elasticity of a cloud-based storage, and, then, present the design and evaluation of ElastMan, an elasticity controller for Cloud-based elastic key-value stores. ElastMan combines feedforward and feedback control. Feedforward control is used to respond to spikes in the workload by quickly resizing the service to meet SLOs at a minimal cost. Feedback control is used to correct modeling errors and to handle diurnal workload. To address nonlinearities, our design of ElastMan leverages the near-linear scalability of elastic Cloud services in order to build a scale-independent model of the service. We have implemented and evaluated ElastMan using the Voldemort key-value store running in an OpenStack Cloud environment.

Our evaluation shows the feasibility and effectiveness of our approach to automation of Cloud service elasticity.





## Motivation

[www.kth.se](http://www.kth.se)

### Elasticity Manager for Elastic Key-Value Stores in the Cloud

Vladimir Vlassov [vlad.v@kth.se](mailto:vlad.v@kth.se)  
KTH Royal Institute of Technology

Stockholm, Sweden

Joint work with: Ahmad Al-Shishtawy, SICS

Cloud Control Workshop, Lund University, 7-9 May 2014



- Web 2.0 applications
  - Wikis, social networks, media distribution and sharing
  - Data-intensive applications; big data
- Challenges
  - **scalability, elasticity**
    - \* Rapidly growing number of users and amount of user-generated data, data-intensive applications
  - **load balancing, latency**
    - \* Uneven load; users are geographically scattered
  - **availability**
    - \* Partial failures, very high load, load spikes
  - **consistency guarantees**
    - \* Data-centric applications and services

May 2014, Lund

Vladimir Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud



### Cloud-Based Services and Applications

2

- **Clouds** provide the illusion of the **infinite amount of resources**
  - **Pay-as-you-go** pricing model
    - High load: Allocate more resources to improve performance
    - Low load: Release resources to save money
  - Enables **Cloud-based Elastic Services and Applications**

Vladimir Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud



### The Need for Elasticity

3

- Web services and applications frequently experience **high workloads**
  - A service can become **popular** in just an hour
- The high level load does not last for long and keeping resources in the Cloud **costs money**
- This has led to **Elastic Computing**
  - Ability of a system to grow and shrink at run-time in response to changes in workload
  - **Cloud computing** allows on-the-fly requesting and releasing VMs to scale/resize the service **in order to meet SLOs at a minimal cost**

Vladimir Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

4

Vladimir Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

May 2014, Lund

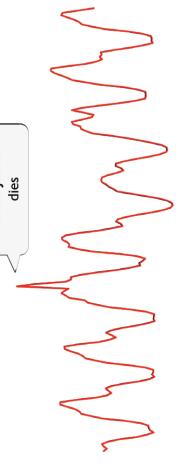
## Wikipedia workload trace - June 2009



- In Physics, **Elasticity** is “the property of a body or substance that enables it to resume its original shape or size when a distorting force is removed” [The Free Dictionary]

- In Cloud computing, **Elasticity** is the ability to scale resource usage up and down [rapidly] according to [instantaneous] demand
  - The ability of a system to scale up and down (grow and shrink by requesting and releasing resources) in response to changes in its environment, workload, and QoS requirements

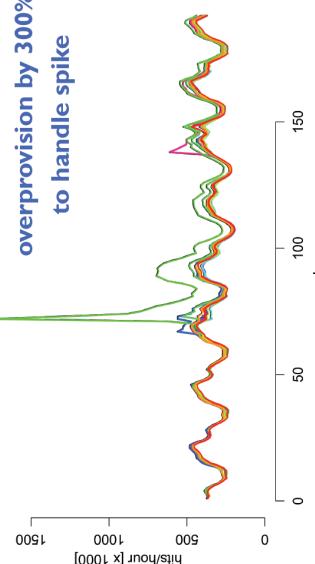
Michael Jackson  
dies



**overprovision by 25%**  
**to handle spike**

## Over-provisioning a Storage System

**over-provision by 300%**  
**to handle spike**



(assuming data stored on ten servers)  
[From: “Solving the Scalability Dilemma with Clouds, Crowds, and Algorithms”, Michael Franklin, UC Berkeley ]

7

5

May 2014, Lund  
Vladimir Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

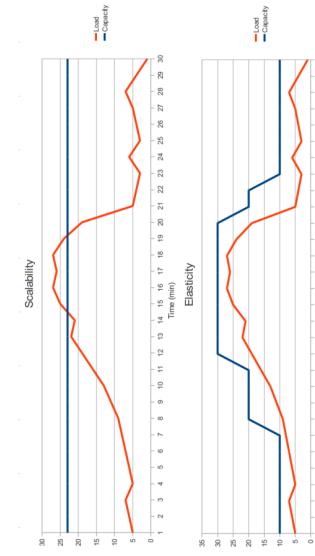
May 2014, Lund  
Vladimir Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

6

May 2014, Lund  
Vladimir Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

6

## Elasticity versus Static Provisioning



Vladimir Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

8

7

7

7

7

7



BY

ND

NC

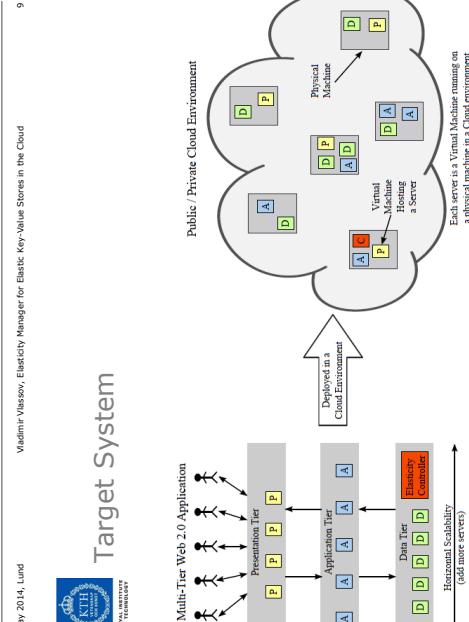


## Outline

- Automation of elasticity
  - Elasticity control for a cloud-based elastic storage
    - Requirements; challenges; approaches
    - Feedforward control to handle workload spikes
    - Feedback control to handle gradual workload changes
  - Evaluation of ElastMan in Voidemont key-value store
  - Conclusions

## Automation of Elasticity

- Elasticity can be controlled either **manually** (by the sys-admin) or **automatically** (by a autonomic manager)
- Automation of elasticity can be achieved by providing an **Elasticity Controller**
  - Helps to avoid SLO violations while keeping the cost low
    - **Automatically adds/removes VMs** (servers, service instances) in response to **changes in some SLO metrics**, e.g., access latency, caused by **changes in workload**
- Can be built using elements of **Control Theory** and/or **Machine Learning** techniques
  - Feedback-loop (a.k.a. closed-loop) control
    - Model Predictive Control (MPC)



May 2014, Lund

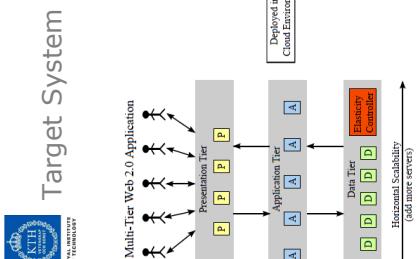
Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

9

Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

10

Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud



May 2014, Lund

Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

11

Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

12

Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

## Storage Services. Key-Value Stores

- **Storage systems designed for horizontal scalability**, such as **key-value stores**
  - minimum functionality: **get(key)** and **put(key, value)**
  - horizontal scalability, load balancing and replication
- Examples
  - Yahoo! PNUTS
  - Google BigTable
  - LinkedIn Voldemort
  - Apache Cassandra
  - UCB's SCADS
  - File systems, e.g., Hadoop Distributed File System



## Storage Services. Key-Value Stores

## An Elasticity Controller for Cloud-Based Key-Value Stores



- Objective
  - To meet **SLOs** at a minimal cost by controlling the **elasticity** (size) of Cloud-based **key-value stores** by adding/removing resources (VMs, storage nodes)
- SLO Examples
  - Average read latency in one minute interval is less than 10ms
  - 99% of reads in one minute interval are performed in less than 10ms per read

May 2014, Lund Vladimír Vlăsov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

13 May 2014, Lund Vladimír Vlăsov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

14

## Requirements for Elastic Storage



- **Horizontal scalability**
  - The storage and I/O capacity **scales (roughly linearly)** with the number of the active servers.
- **Touch points**
  - Sensors to monitor workload and performance (e.g., read latency)
  - Actuators to add/remove resources and service instances
- **Load balancing (re-balancing)**
  - Distribute data across servers to effectively balance load;
  - Redistribute (rebalance) data in response to join/leave events
- **Replication**
  - For robust availability: enough to avoid interruptions on leave events

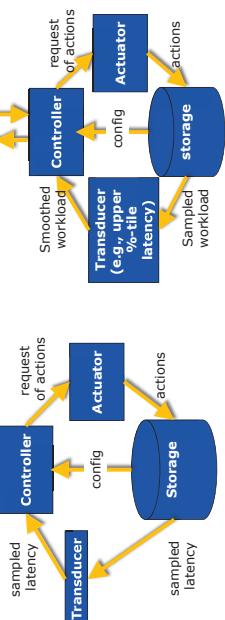
May 2014, Lund Vladimír Vlăsov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

15

## Approaches to Automated Elasticity Control for a Cloud-Based Storage



- Closed-loop control e.g., [Lim2010][Mou2012]
- Model Predictive Control e.g., [Tru2011]



May 2014, Lund Vladimír Vlăsov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

14

## Challenges of Elasticity Control for a Cloud-Based Storage



- **Actuator delays** due to data movement (rebalancing)
- **Interference** with applications and sensor measurements
- **Discrete storage units**
- **Nonlinearity** due to diminishing reward of adding a storage unit with increasing scale
  - 99th percentile of access latency is a relatively **noisy signal**
- VM performance is **difficult to model and predict**
  - Highly dynamic workload that is composed of both **gradual (diurnal) and sudden (spikes) variations**

May 2014, Lund Vladimír Vlăsov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

16



## Touch-points



## Sensors (1/3)

### Sensors

- To monitor **workload** and **performance** (e.g., access latency)
  - A controller input includes one or more system performance metrics (system outputs)
- Requirements to system metrics [Lim 2010]
  - Easy** to measure accurately
  - Should expose the **system(ter)-level behavior** or performance
  - Should be reasonably **stable**
  - Should **correlate** to the measure of service level specified in SLO

### Actuators

- To add/remove resources and service instances
  - Cloud APIs to request/release server instances (VMs)
  - Storage API to request handling joins and leaves and rebalancing



## Sensors (2/3)

- In many cases, a system performance metric **strongly correlates** with the overall request latency (response time)
- Performance counters** (CPU/memory/network utilizations), can be used as sensors

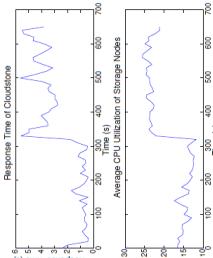


Figure 2: Cloudstone response time and average CPU utilization of the storage nodes, under a light load and a heavy load that is bottlenecked in the storage tier. CPU utilization in the storage tier correlates strongly with overall response time (the coefficient is .88), and is a more stable feedback signal. [Lim2010]

May 2014, Lund Vladimír Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

May 2014, Lund Vladimír Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud



## Sensors (3/3)

### SLO requirements (what to regulate)

- Average** values (means) of the SLO metrics
  - E.g., an average response time, an average CPU utilization
  - Classical closed-loop control
  - HSC integral controller of HDFS-based storage [Lim2010]
- Upper quantiles** of the SLO metrics
  - E.g. 99<sup>th</sup> percentile of latency
  - 99% of all requests must be answered within 100ms<sup>“</sup>
  - Model Predictive Control
  - Example: SCADS Director for SCADS storage (UCB) [Tru2011]

May 2014, Lund Vladimír Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

May 2014, Lund Vladimír Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

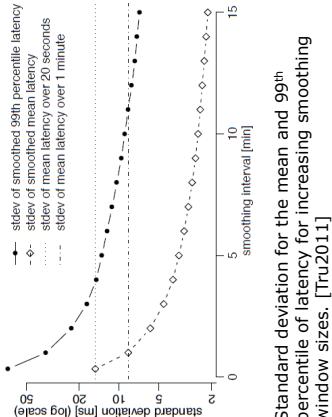
May 2014, Lund Vladimír Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

## ElastMan: Elasticity Manager for Cloud-Based Key-Value Stores



### Mean versus Upper Quartile

- Usually, upper percentiles of latency measurements are very “noisy”
- A noisy latency signal can cause oscillations in classical closed-loop control



Standard deviation for the mean and 99<sup>th</sup> percentile of latency for increasing smoothing window sizes. [Tru2011]

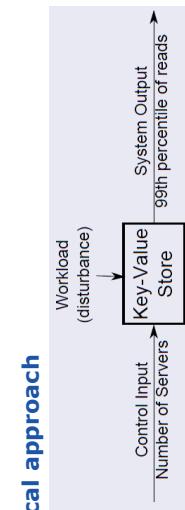
May 2014, Lund Vladimír Vlăsov, Elasticity Manager for Elastic Key-Value Stores in the Cloud 21

22

### Modeling the Store (1/2)



### Typical approach



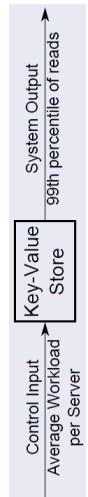
#### Non linear Model

- Adding 1 node to a 1 node system -> doubles the performance
- Adding 1 node to a 100 nodes system -> only 1% improvement
- Workload treated as **disturbance**

May 2014, Lund Vladimír Vlăsov, Elasticity Manager for Elastic Key-Value Stores in the Cloud 23

24

### ElastMan approach



- Control the number of servers **indirectly** by controlling the **average workload per server**
- Relyes on **near linear scalability** of key-value stores

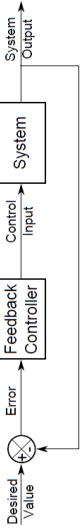
22

Vladimír Vlăsov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

23

Vladimír Vlăsov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

## Feedback Control [Hel2004]



- The system's **output** (e.g., response time) is being **monitored**
- Controller calculates the **control error**
- Controller changes the control input (e.g., number of servers to add or remove) according to the **amount** and **sign** of the control error
- Advantage:** controller can adapt to **disturbance**
- Disadvantages:** oscillation, overshoot, possible instability

May 2014, Lund Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

25 May 2014, Lund Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

## Feedforward Control [Hel2004]

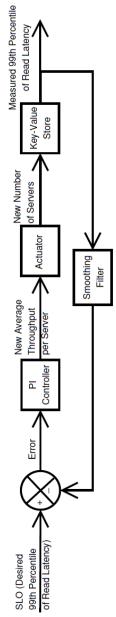


- The system's **output** is **not monitored**
- Other system states and variables are **monitored**
- Controller relies on a **model of the system** to calculate necessary change
- Advantages:** faster and avoids oscillation and overshoot
- Disadvantages:** **sensitive to unexpected disturbances** that are not modeled

26 May 2014, Lund Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

26 May 2014, Lund Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

## ElastMan Feedforward Controller

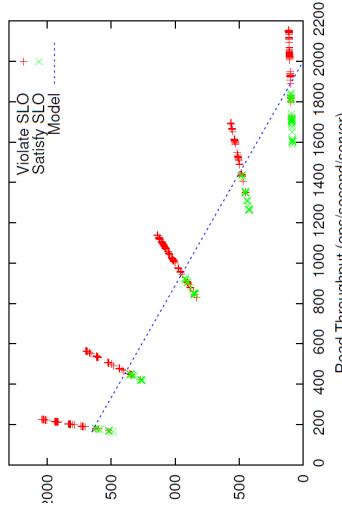


May 2014, Lund Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

27 May 2014, Lund Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

28 May 2014, Lund Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

## Binary Classifier



Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

29

## Feedback

- Classical PI controller
- Monitors 99th percentile of read latency
  - E.g., the read latency of 99% of reads in a 1 min interval is at most 10ms
- Can tolerate and adapt to **modeling errors**
- Allows **smoothing** the noisy 99th percentile signal

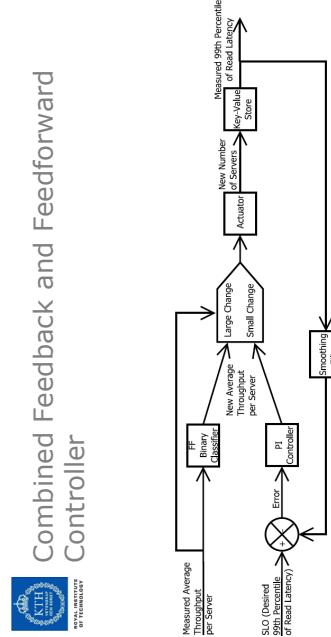
30

## Feedforward

- A **binary classifier** using logistic regression
- Deals with workload **spikes** (large rapid changes)
- Monitors 99th percentile of read latency
- Monitors workload (intensity of reads and intensity of writes)

31

## Combined Feedback and Feedforward Controller



May 2014, Lund

32



## ElastMan: Combining Feedback and Feedforward Control

## Feedforward

- A **binary classifier** using logistic regression
- Deals with workload **spikes** (large rapid changes)
- Monitors 99th percentile of read latency
- Monitors workload (intensity of reads and intensity of writes)

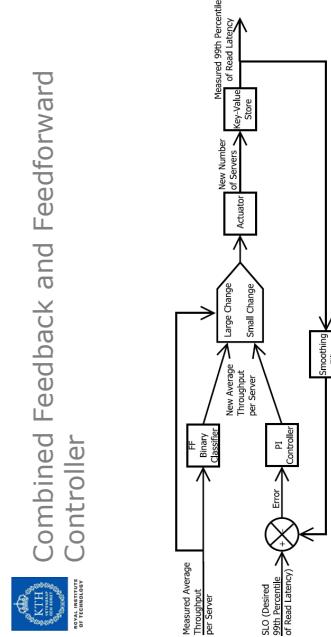
33

## Feedback

- Classical PI controller
- Monitors 99th percentile of read latency
  - E.g., the read latency of 99% of reads in a 1 min interval is at most 10ms
- Can tolerate and adapt to **modeling errors**
- Allows **smoothing** the noisy 99th percentile signal

34

## Combined Feedback and Feedforward Controller



May 2014, Lund

35

Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

31

Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

32

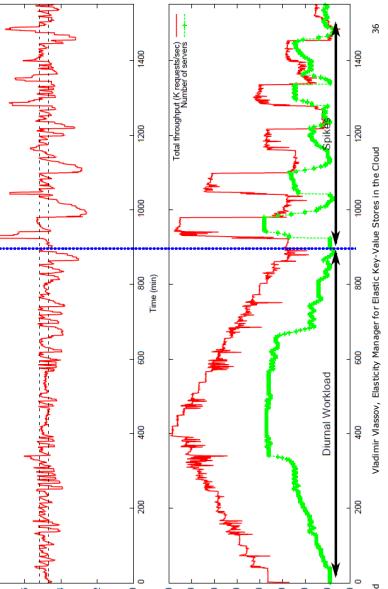
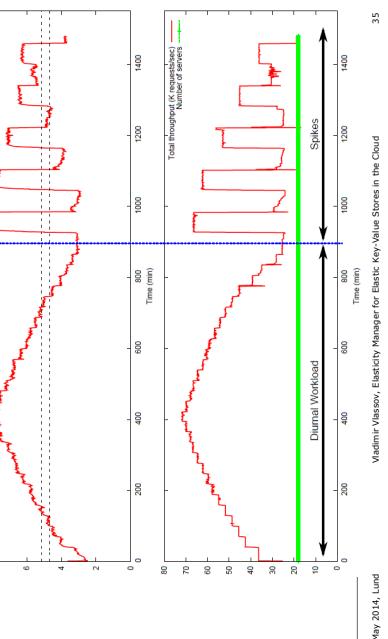
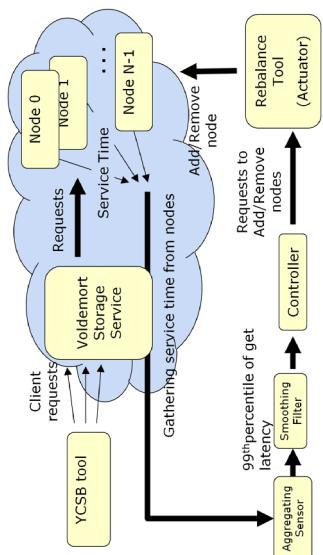


## Evaluation



### Voldemort Key-Value Store with the ElastMan Elasticity Controller

- Implemented a prototype of the **ElastMan** elasticity controller
- Evaluated with LinkedIn's Voldemort key-value store
- Deployed in our private OpenStack Cloud



33

May 2014, Lund

Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

May 2014, Lund

Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

34

May 2014, Lund

Vladimir Vassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud

## Conclusions



## ElastMan References

- ElastMan addresses the challenges of the variable performance of Cloud VMs, dynamic workload, and stringent performance requirements
- ElastMan combines and leverages the advantages of both feedback and feedforward control
  - feedforward control quickly responds to rapid changes in workload
  - feedback controller handles diurnal workload and to correct modeling errors in the feedforward control
- Evaluation results show the feasibility of the ElastMan approach



## References

- Ahmad Al-Shishitawy, Vladimir Vlassov, **ElastMan: Elasticity Manager for Elastic Key-Value Stores in the Cloud**, The ACM Cloud and Autonomic Computing Conference (CAC 2013), Miami, FL, USA, August 5–9, 2013.
- Ahmad Al-Shishitawy, Vladimir Vlassov, **ElastMan: Autonomic Elasticity Manager for Cloud-Based Key-value Stores**, The 22nd ACM International Symposium on High-Performance Parallel and Distributed Computing (HPDC '13), ACM, New York, NY, USA, pp. 115-116.
- <http://www.ict.kth.se/~ahmadas/ElastMan/>

May 2014, Lund	Vladimír Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud	38
May 2014, Lund	Vladimír Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud	37

- [Hel2004] J. L. Hellerstein, Y. Diao, S. Parekh, and D. M. Tilbury, *Feedback Control of Computing Systems*. Wiley-IEEE Press, 2004.
- [Lim2010] Harold C. Lim, Shivnath Babu, and Jeffrey S. Chase, *Automated control for elastic storage*. ICAC '10, 2010
- [Mou2012] M. A. Moulaei, A. Al-Shishitawy, and V. Vlassov, *State-Space Feedback Control for Elastic Distributed Storage in a Cloud Environment*, ICAS 2012
- [Tru2011] Beth Trushkowsky, et al. *The SCADS director: scaling a distributed storage system under stringent performance requirements*. The 9th USENIX Conf on File and Storage Technologies (FAST'11), 2011



- May 2014, Lund Vladimír Vlassov, Elasticity Manager for Elastic Key-Value Stores in the Cloud 39

**PRINCIPLES AND METHODS FOR ELASTIC COMPUTING**  
**Schahram Dustdar, Vienna University of Technology**

Elasticity is seen as one of the main characteristics of Cloud Computing today. Is elasticity simply scalability on steroids? In this talk I will discuss the main principles of elasticity, present a fresh look at this problem, and examine how to integrate people, software services, and things into one composite system, which can be modeled, programmed, and deployed on a large scale in an elastic way.



## Acknowledgements

# Principles and Methods for Elastic Computing -

Lund, 7 May 2014

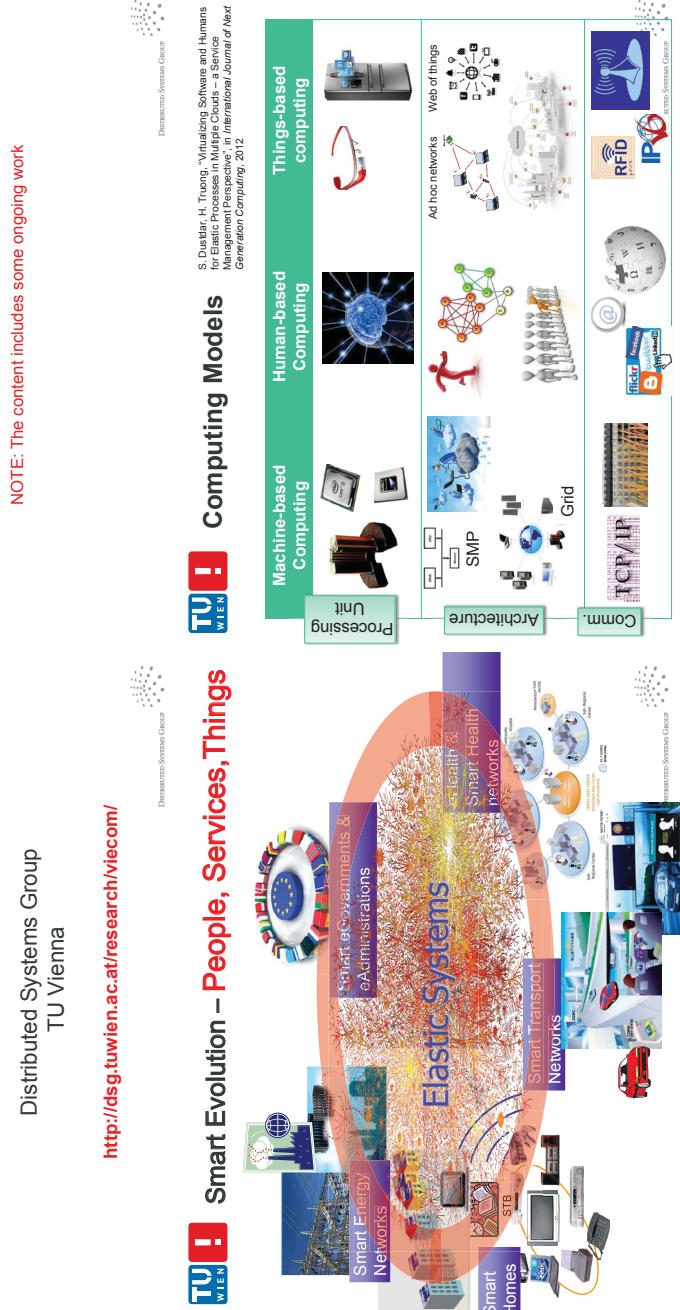
Schahram Dustdar  
Distributed Systems Group  
TU Vienna

<http://dsg.tuwien.ac.at/research/viecom/>

Includes some joint work with Hong-Linh Truong, Alessio Gambi, Muhammad Z. C. Candia, Georgiana Copil, Duc-Hung Le, Daniel Moldovan, Stefan Nasic, Mirela Riveri, Sanjin Sehic, Ognjen Sasic



NOTE: The content includes some ongoing work

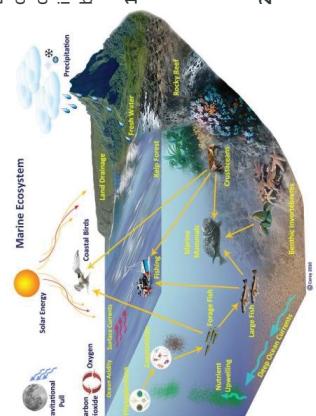


## TU WIEN Think Ecosystems: People, Services, Things



### Our Approach

- **Unified service unit model** (Consumption, ownership, provisioning, price, function, etc.)
- **Connecting Data Centers to IoT**
  - From physically isolated verticals to virtual verticals
  - **Software-defined elastic data centers and IoT ecosystems**
  - SD units are described with well-defined API
  - Provisioning units for customized gateways
  - Dynamically composing units into runtime topologies
  - Runtime controlling and optimization via configuration policies (DevOps principle)
- **Human Augmentation**
  - Human computation capabilities under elastic service units
  - Programming human-based units together with software-based units



## TU WIEN Elasticity ≠ Scalability



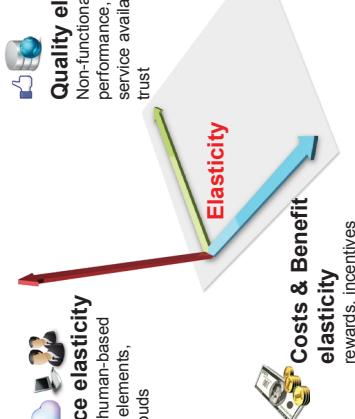
### Quality elasticity

Non-functional parameters e.g., performance, quality of data, service availability, human trust



### Resource elasticity

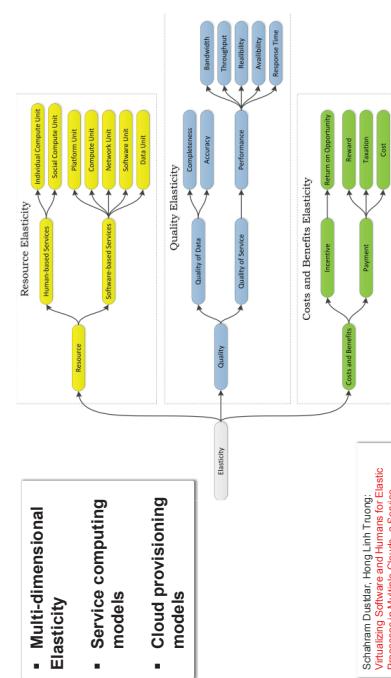
Software / human-based computing elements, multiple clouds



Dienstleistungs Systeme Gruppe

Dienstleistungs Systeme Gruppe

## TU WIEN Vienna Elastic Computing Model



Dienstleistungs Systeme Gruppe

Dienstleistungs Systeme Gruppe

Schäfman Dietmar, Hong Lam Truong; Virtualizing Software and Humans for Elastic Processes in Multiple Clouds - a Service Management Perspective, UNGC 3(2) (2012)

# TU Wien Diverse types of elasticity requirements

- Application user:** "If the cost is greater than 800 Euro, there should be a scale-in action for keeping costs in acceptable limits"
- Software provider:** "Response time should be less than amount X varying with the number of users."
- Developer:** "The result from the data analytics algorithm must reach a certain data accuracy under a cost constraint. I don't care about how many resources should be used for executing this code."

**Cloud provider:** "When availability is higher than 99% for a period of time, and the costs is the same as for availability 80%, the cost should increase with 10%."

## TU Wien High Level Description of Elasticity Requirements

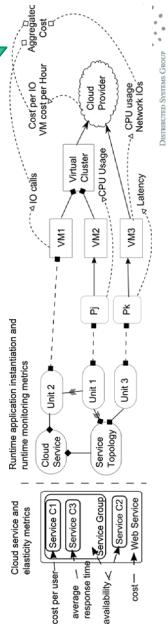
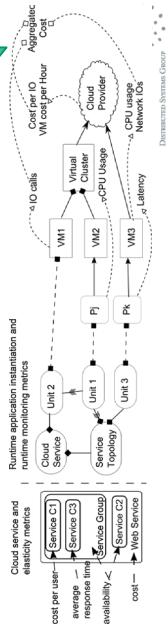
SYBL (Simple Yet Beautiful Language) for specifying elasticity requirements  
SYBL-supported requirement levels

1. Cloud Service Level
2. Service Topology Level
3. Service Unit Level
4. Relationship Level
5. Programming/Code Level

```
#SYBL:CloudServiceLevel
Cons1: CONSTRAINT responseTime < 5 ms
Cons2: CONSTRAINT responseTime < 10 ms
WHEN inDutliers > 10000
Str1: STRATEGY CASE fulfilled(Cons1) OR
fulfilled(Cons2); minimize (cost)

#SYBL:ServiceUnitLevel
Str2: STRATEGY CASE (oCost < 3 Euro ;
maximize( dataFreshness ) )

#SYBL:CodeRegionLevel
Cons4: CONSTRAINT dataAccuracy > 90%
AND cost < 4 Euro
```



## TU Wien Data Center - Engineering Techniques

### TU Wien Mapping Services Structures to Elasticity Metrics



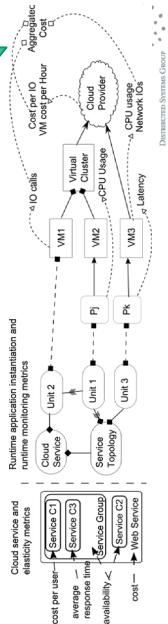
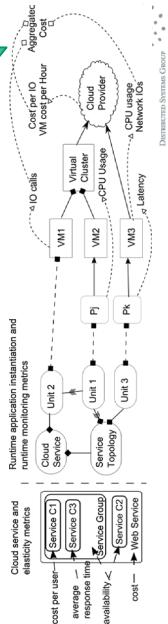
## TU Wien

### TU Wien High Level Description of Elasticity Requirements

```
#SYBL:CloudServiceLevel
Cons1: CONSTRAINT responseTime < 5 ms
Cons2: CONSTRAINT responseTime < 10 ms
WHEN inDutliers > 10000
Str1: STRATEGY CASE fulfilled(Cons1) OR
fulfilled(Cons2); minimize (cost)

#SYBL:ServiceUnitLevel
Str2: STRATEGY CASE (oCost < 3 Euro ;
maximize( dataFreshness ) )

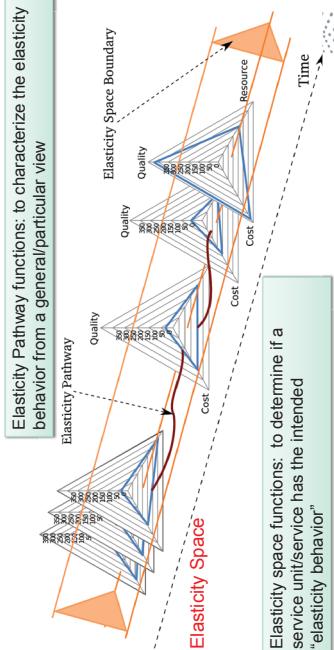
#SYBL:CodeRegionLevel
Cons4: CONSTRAINT dataAccuracy > 90%
AND cost < 4 Euro
```



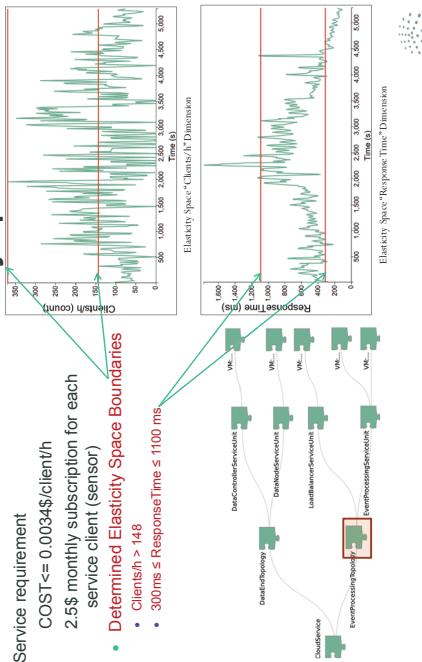
Georgiana Copil, Daniel Moldovan, Hong-Linh Toong, Scaphram Duadar, **"SYBL: an Extensible Language for Controlling Elasticity in Cloud Applications"**, 13th IEEE/ACM International Symposium on Cluster, Cloud and Grid Computing (CCGrid), May 14-16, 2013, Delft, Netherlands

## TU WEN ! Elasticity Model for Cloud Services

Moldovan D., G. Copil, Truong H.-L., Dusitdar S. (2013). **MELA: Monitoring and Analyzing Elasticity of Cloud Service**. CloudCom 2013



## TU WEN ! Multi-Level Elasticity Space

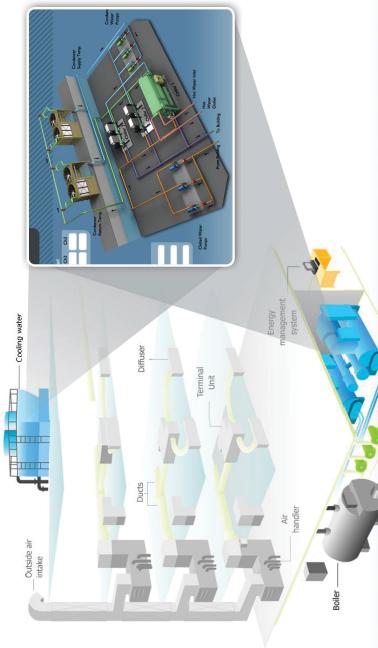


## IoT-Engineering Techniques

### TU WEN !

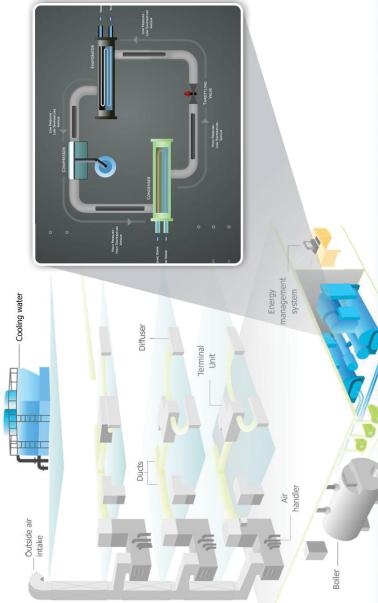
Dimension Summary Grid

### Water Ecosystem



© Copyright 2010 Pacific Control Systems. All Rights Reserved.

### HVAC (Heating, Ventilation, Air Conditioning) Ecosystem



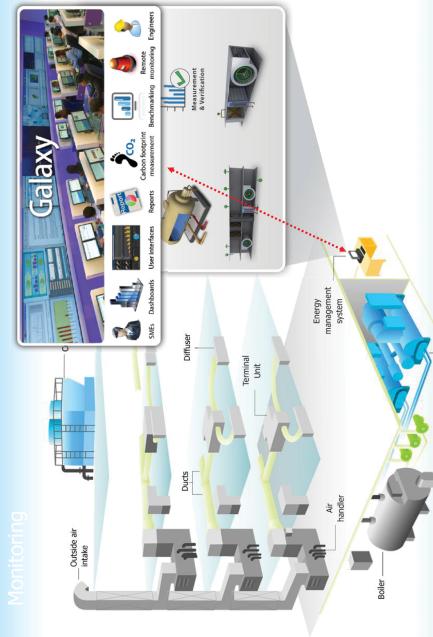
© Copyright 2010 Pacific Control Systems. All Rights Reserved.

### Air Ecosystem



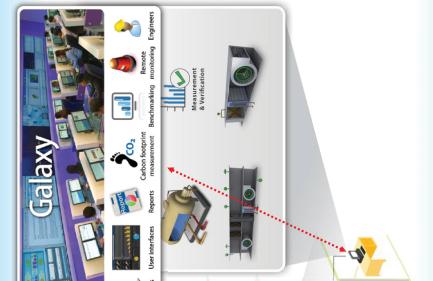
© Copyright 2010 Pacific Control Systems. All Rights Reserved.

### Monitoring



© Copyright 2010 Pacific Control Systems. All Rights Reserved.

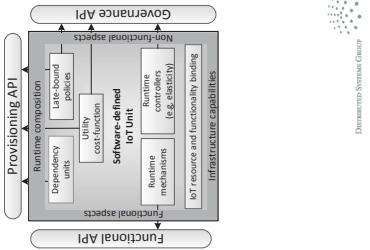
### Water Ecosystem



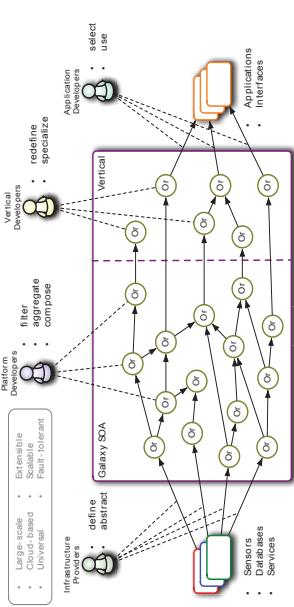
© Copyright 2010 Pacific Control Systems. All Rights Reserved.

## TU WIE ! Software-defined IoT units

- Provide **software-defined API** for accessing, configuring and controlling units
- Support fine-grained internal **configurations**, e.g. adding functional capabilities like different communication protocols, at runtime.
- Can be **composed** at higher-level, via dependency units, creating **virtual topologies** (of multiple gateways) that can be (re)configured at runtime.
- Enable decoupled and managed configuration (via late-bound policies) to **provision the units dynamically** and on-demand.
- Have utility cost-functions that enable pricing the IoT resources as utilities.



## TU WIE ! The Programming Model (Origins)



Sehic, S., Li, F., Nastic, S., Dusdar, S.: A Programming Model for Context-Aware Applications in Large-Scale Pervasive Systems, The 8th IEEE International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob 2012), 8-10 October 2012, Barcelona, Spain.

Sehic, S., Li, F., Nastic, S., Vringer, M., Li, F., Dusdar, S.: Entity-Adaptation - A Programming Model for Dynamic Adaptation of Context-Aware Applications, 2014 Symposium on Applied Computing (SAC 2014), Mobile Computing and Applications (MCA) Track, 24-28 March 2014, Gyeongju, Republic of Korea

**Chiller Plant Analysis Tool**

Chiller Performance Metrics

Parameter	Value
41°C	Outside Air Temperature
21.8%	Refrigerant Load
65.5 kW	Electric Load
131.4 A kWh	Efficiency
100% / 100%	Coil Return / coil Outflow

Compressor A

MOTOR CURRENT: 1000 A

MOTOR CURRENT: 665.4 A

MOTOR TEMPERATURE: 99.0 °C

DISCHARGE GAS TEMPERATURE: 46.7 °C

DISCHARGE GAS PRESSURE: 11.76 psi

SUCTION PRESSURE: 4.40 psi

SATURATED SUCTION TEMPERATURE: 2.8 °C

Oil Pressure: 100.9 psi

Oil Pressure: 51.4 psi

Difference: 49.5 psi

Saturated Condensing Temperature: 10.2 °C

Oil Pressure: 45.9 psi

To Compressor: 21.6 °C

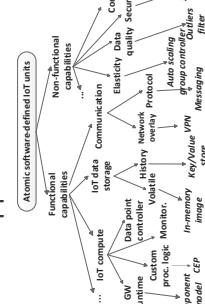
Oil Pressure: 2.4 psi

Difference: 26.1 °C

Saturated Condensing Temperature: 26.1 °C

## TU WIE ! Ecosystem for software-defined IoT systems

- Create an ecosystem of software-defined IoT units for the creation of software-defined IoT systems.
- Distributing IoT units in a market-like fashion, e.g., via IoT AppStore.

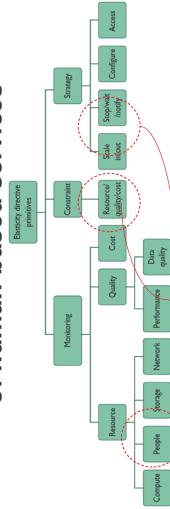


Darmstadt Systems Group

Darmstadt Systems Group

## TU! Specifying and controlling elasticity of human-based services

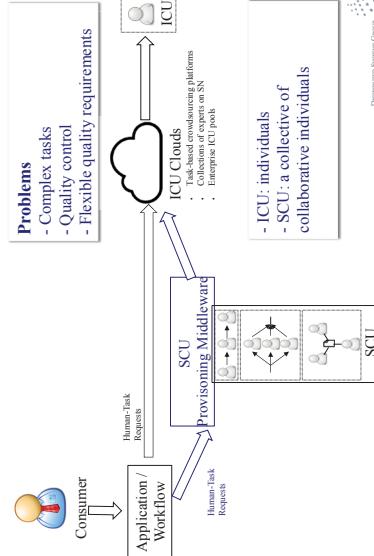
### Human augmentation – Engineering Techniques



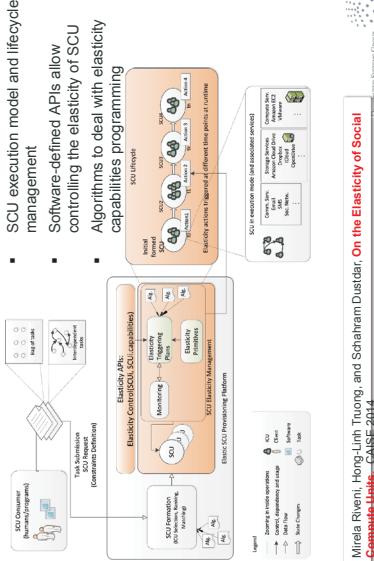
What if we need to invoke a human?

## for service unit analyzing chiller measurement  
#SYBL!\_ServiceUnit\_Level  
Mon1 MONITORING accuracy = Quality Accuracy  
Const1 CONSTRAINT accuracy < 0.7  
Const1 CONSTRAIN accuracy < 0.7  
SRI1 STRATEGY CASE Violated(Cons1):  
NotifyIncident(DEFAULT\_ServiceUnitType,HBS)

## TU! scu for independent tasks



## TU! Elasticity Capabilities and APIs



Mirela Riveni, Hong-Linh Tuong, and Sandrahm Dustdar, **On the Elasticity of Social Compute Units**, CAiSE 2014

Dimitrios Syntetos-Giannakouris, **Intelligent Systems for the Internet of Things**

## **TU !** Conclusions (1) – Engineering Elasticity

- The evolution of underlying systems and the utilization of different types of resources under different models for elasticity requires
- Complex, open **hybrid service unit provisioning frameworks**
- Different **strategies** for dealing with different types of tasks
- **Quality issues** for software, data, and people in an integrated manner for different perspectives
- We are just at an early stage of developing techniques for engineering elastic applications wrt multi-dimensional elasticity



## **TU !** Conclusions (2) – Engineering Elasticity

- Service engineering analytics of elastic systems**
  - Programming hybrid compute units for elastic processes
  - Elasticity specifications and reasoning techniques
  - Elasticity spaces analytics
- Application domains**
  - “Social computer” and smart cities (FP 7 FET Smart Cities and PC3L)
  - Computational science and engineering (FP 7 CELAR)



Thanks for your attention!



Prof. Dr. Schahram Dustdar  
Distributed Systems Group  
TU Wien

[dsg.tuwien.ac.at](http://dsg.tuwien.ac.at)



**EXPLORING AUTONOMICS FOR CLOUDS****Manish Parashar, Rutgers University**

A grand challenge for verification of control systems could be stated as follows: Given requirement Cloud computing has emerged as a dominant paradigm that is being widely adopted by enterprises. Clouds and Cloud federations are also rapidly joining high-performance computing system, clusters and Grids as viable platforms for scientific exploration and discovery. Clouds offer on-demand access to computing utilities, an abstraction of unlimited computing resources, customizable environments, and a pay-as-you-go business model. They also provide a potential for dynamic scale-up, scale-down and scale-out, and support IT outsourcing and automation. As a result, it is possible to create hybrid federated cloud infrastructures integrating private clouds, local data centers, and public clouds. However, developing and managing cloud applications/services to appropriately use the capacities and capabilities offered by these cloud federations can be challenging – for example, applications/services need to be managed according to pricing policy, quality of service requirements, budgets, etc. In this talk, I will explore autonomic application execution and management in federated Cloud infrastructures.



## Moving towards the Cloud

### Exploring Autonomics for Federated Clouds

Moustafa AbdelBaky, Javier Diaz-Montes, Mengsong Zou and **Manish Parashar**

The NSF Cloud and Autonomic Computing Center  
 Rutgers, The State University of New Jersey, USA  
<http://cometcloud.org>

In collaboration with Omer Rana, Tom Beach, Ioan Petri, Cardiff University, UK



- Cloud services provide an attractive platform for supporting the computational and data needs of academic and business application workflows

- Cloud paradigm:

- Rent resources as cloud services on-demand and pay for what you use
  - Potential for scaling-up, scaling-down and scaling-out, as well as for IT outsourcing and automation
- Hybrid cloud services landscape spanning private clouds, public clouds, HEC centers, etc.
  - Heterogeneous offering with different QoS, pricing models, availability, capabilities, and capacities

## Cloud Federations – Motivations

- Application workflow exhibit heterogeneous and dynamic workloads, and highly dynamic demands for resources
  - Various and dynamic QoS requirements
    - Throughput, budget, time
  - Often involve large amounts of data
    - Large size, heterogeneous nature, and geographic location
- Such workloads are hard to be efficiently supported using classic federation models

- Implications of the cloud paradigm
  - Rent required resources as cloud services on-demand and pay for what you use
    - Heterogeneous offering with different QoS and costs
- Provisioning and federating an appropriate mix of resources on-the-fly is essential and non-trivial

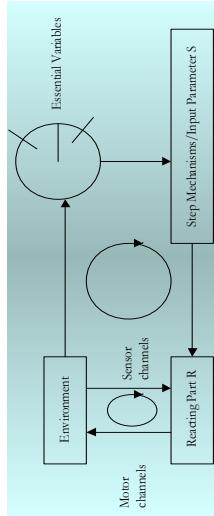
## AUTONOMICS FOR CLOUD FEDERATIONS

## RUTGERS

### Integrating Biology and Information Technology: The Autonomic Computing Metaphor (~2004)

- Current paradigms, mechanisms, management tools are inadequate to handle the scale, complexity, dynamism and heterogeneity of emerging systems and applications
- Nature has evolved to cope with scale, complexity, heterogeneity, dynamism and unpredictability, lack of guarantees
  - self configuring, self adapting, self optimizing, self healing, self protecting, highly decentralized, heterogeneous architectures that work !!!
- Goal of autonomic computing is to enable self-managing systems/applications that addresses these challenges using high level guidance
  - Separation of policy and mechanisms; Holistic; Automation

*"Autonomic Computing: An Overview," M. Panaster, and S. Haifi, Hot Topics, Lecture Notes in Computer Science, Springer Verlag, Vol. 3566, pp. 247-259, 2005.*



### Ashby's Ultrastable System (1920s)

## RUTGERS

### Integrating Biology and Information Technology: The Autonomic Computing Metaphor (~2004)

- Rich body of work on using autonomics for cloud/data-center management
  - Provisioning
  - Workload management
  - Power/energy management
  - Etc. . .
  - Using control theoretic approaches



### Autonomic Cloud/ACI Federation

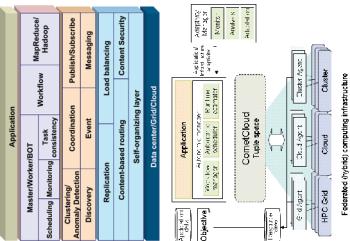
- Assemble a federated cloud/ACI on-the-fly integrating clouds, grids and HPC
  - Cloud-bursting, dynamic application scale-out/up to address dynamic workloads, spikes in demand, and other extreme requirements
  - Cloud-bridging, on-the-fly integration of different resource classes
- Provide policy-driven autonomic resource provisioning, scheduling and runtime adaptations
  - What and where to provision?
  - Policies encapsulate user's requirements (deadline, budget, etc.), resource constraints (failure, network, availability, etc.)
- Provide programming abstractions to support application workflows

*"Autonomic Computing: An Overview," M. Panaster, and S. Haifi, Hot Topics, Lecture Notes in Computer Science, Springer Verlag, Vol. 3566, pp. 247-259, 2005.*

## CometCloud – Federated Clouds for Science

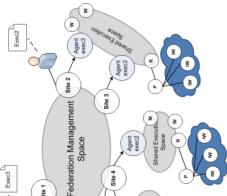
- Enable applications on dynamically federated, hybrid infrastructure exposed using Cloud abstractions
  - Services; discovery, associative object store, messaging, coordination
  - Cloud-bursting; dynamic application workloads, spikes in demand, and extreme requirements
    - Cloud bridging on-the-fly migration of different resource classes (public & private clouds, data-centers and HPC Grids)
- High-level programming abstractions & automatic mechanisms
  - Cross-layer Autonomics; Application layer, Service layer, Infrastructure layer
- Diverse applications
  - Business intelligence, financial analytics, oil reservoir simulations, medical informatics, document management, etc.

<http://cometcloud.org>



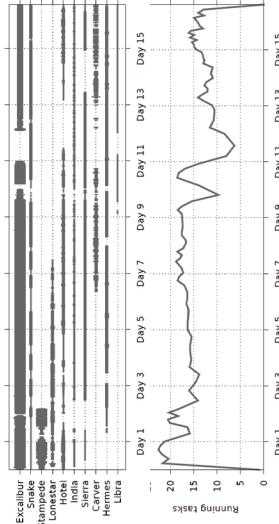
## On-Demand Elastic Federation using CometCloud

- Software defined ACI federations exposed using elastic on-demand Cloud abstractions
- Autonomic cross-layer federation management using user and provider policies and constraints
  - Separately defined; dynamically evolving
    - Specified based on availability, cost/ performance constraints, etc.
    - Sites discover/coordinate with each other to:
      - Identify themselves / Verify identity (X.509, public/private key,...)
      - Advertise their own resources, capabilities, availability, constraints
      - Discover available resources
- Federated ACI tested!



## UberCloud Experiment

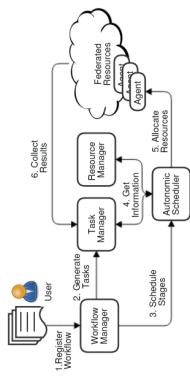
- 10 different resources from 3 countries federated using CometCloud
  - 16 days, 12 hours, 59 minutes and 28 seconds of continuous execution
  - 12,845 tasks processed, 2,897,390 CPU-hours consumed, 400 GB of data generated



## Enabling Data-Driven Workflows

- Enable the autonomic execution of complex workflows in software-defined multi-cloud environments
- Elastically compose appropriate cloud services and capabilities to ensure that the user's objectives are met

# DATA-DRIVEN WORKFLOWS [CLOUD'14] (WITH IBM)



## Optimizing Resource Usage in Multi-Clouds

- Execute a data-driven workflow in a multi-cloud environment
- Different scheduling policies and objectives
  - Minimum Completion Time
  - Centralized storage vs Distributed storage
  - Deadline-based Policy
- Performance optimization (Proc)
  - Data locality optimization (Data)
  - Performance and data optimization (ProcData)
  - Cost optimization (Cost)

## RUTGERS

## Experiment Setup

- Montage workflow
- Three heterogeneous and geographically distributed clouds

VM type <sup>1</sup>	# of Cores	Memory	Max. VMs <sup>2</sup>	Speedup
Alamo - Large	4	8 GB	2	3.55
Alamo - Medium	2	4 GB	4	2.77
Alamo - Small	1	2 GB	2	1.08
Serra - Medium	2	4 GB	2	1.08
Serra - Small	1	2 GB	3	0.71
Hotel - Small	1	2 GB	6	0.76

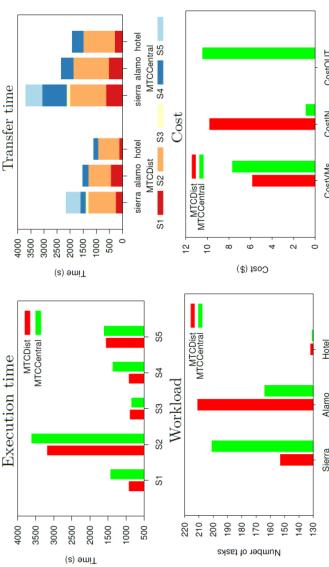
Note: <sup>1</sup> Name of the site followed by the type of VM.  
<sup>2</sup> Maximum number of available VMs per type.

Network (Down/Up)	Alamo	Serra	Hotel
Internal Network	11/2.3	30/30	45/45
External Network	10/0.9	15/1.5	11/1.1
Internet	18/18	12/1	12/1

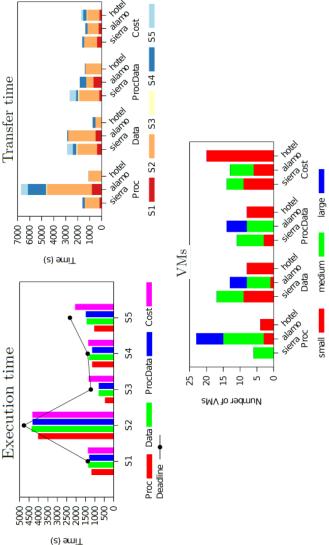


- Hotel – Cloud
- Serra – Cloud
- Alamo – Cloud
- Hotel – U. Chicago

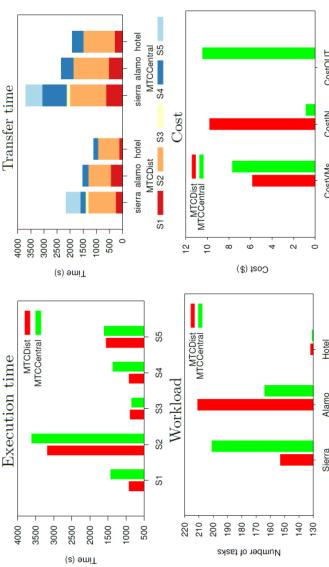
## Minimum Completion Time



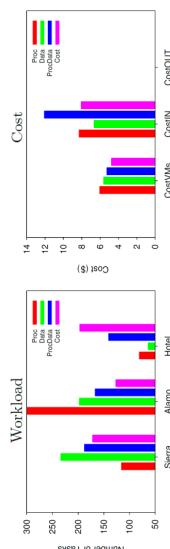
## Deadline-based Policies



## Deadline-based Policies (Cont.)



## Deadline-based Policies

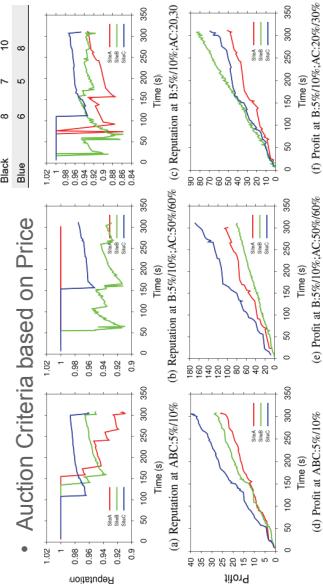


## FEDERATING RESOURCES USING SOCIAL MODELS [IC2E'14]

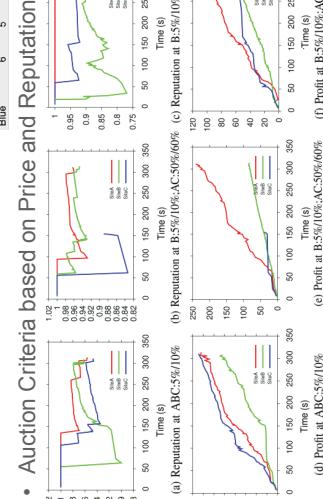
## Exchanging Resources in a Federated Cloud

- Consider federation policies and determine their impact on the overall status of each site
- Market model for resource sharing
  - External task vs Local task
  - Heterogeneous tasks - different deadlines and costs
  - Each site decides how much benefit per task (% cost)
  - Federation policy = Auction criteria
- Federation infrastructure between Cardiff (UK) and Rutgers (USA)

## Profit and Reputation of Each Site

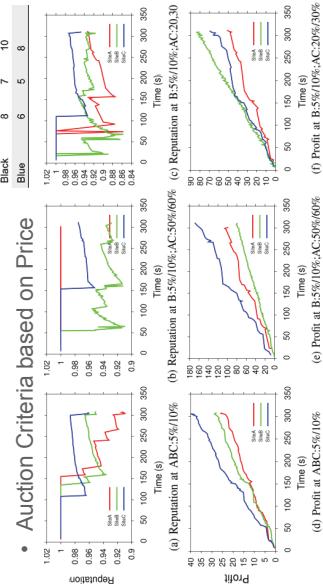


## Profit and Reputation of Each Site II

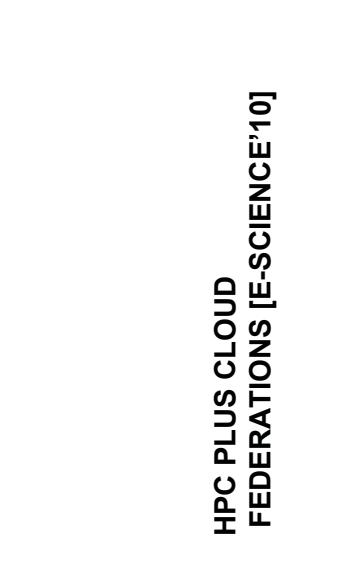


## HPC PLUS CLOUD FEDERATIONS [E-SCIENCE'10]

## Profit and Reputation of Each Site



## Profit and Reputation



## Exploring Hybrid HPC-Grid/Cloud Usage Modes (eScience'09, ScienceCloud'10)

What are appropriate usage modes for hybrid infrastructure?

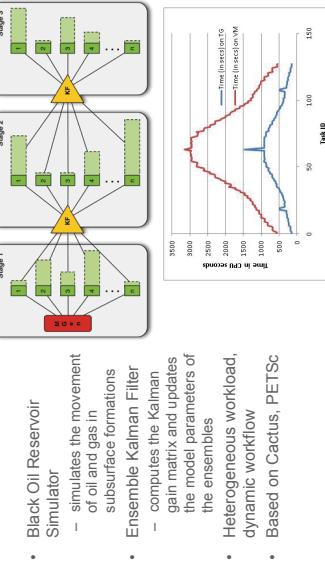
- Acceleration – How can Clouds be used as accelerators to improve the application time to completion
  - To alleviate the impact of queue wait times
  - ‘Strategically’ Off load appropriate tasks to Cloud resources
  - All while respecting budget constraints.

• Conservation – How Clouds can be used to conserve HPC Grid allocations, given appropriate runtime and budget constraints.

• Resilience – How Clouds can be used to handle:

- General Response to dynamic execution environments
- Specific: Unanticipated HPC Grid downtime, inadequate allocations or unexpected Queue delays/OS change

## Reservoir Characterization: EnKF-based History Matching

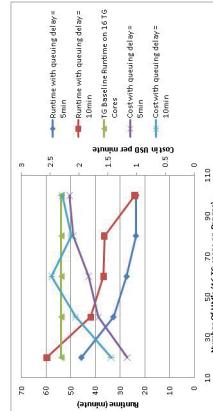


## Using Clouds as Accelerators for HPC Grids

- Explore how Clouds (EC2) can be used as accelerators for HPC Grid (TG) workloads
  - 16 CPUs (Ranger)
    - Average queuing time for Ranger was set to 5 and 10 minutes
    - Number of EC2 VMs (m1.small) from 20 to 100 in steps of 20
      - VM start up time was about 160 seconds

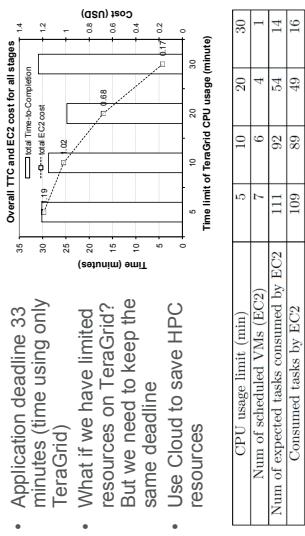
## Using Clouds as Accelerators for HPC Grids I

- Acceleration is more notable with more VMs - lower the TTC
- The reduction in TTC is roughly linear
  - Affected by complex interplay between the tasks in the workload and resource availability



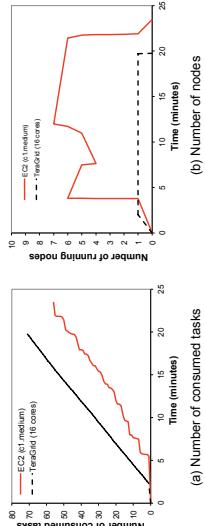
## Exploring Conservation

- Application deadline 33 minutes (time using only TeraGrid)
- What if we have limited resources on TeraGrid? But we need to keep the same deadline
- Use Cloud to save HPC resources



## Exploring Resilience

- Deadline 20 minutes
- Two EC2 instances are failed at around 8 minutes



## Conclusions

- Complex application workflows necessitate software defined federated platforms that integrated heterogeneous cloud services
- Provisioning and federating an appropriate mix of resources on-the-fly is essential and non-trivial
- Autonomics can provide the abstractions and mechanism to manage complexity
  - Separation + Integration + Automation
- However, there are implications
  - Added uncertainty
  - Correctness, predictability, repeatability
  - Variation
  - New formulations necessary...



## The CometCloud Team

**Mousata AfolabiBalogun**  
PhD Student, Dept. of Electrical & Computer Eng.  
Rutgers University  
Email: mafolabi@cac.rutgers.edu

**Mengsong Zou**  
PhD Student, Dept. of Computer Science  
Rutgers University  
Email: mz228@cac.rutgers.edu

**Manish Parashar**, Ph.D.  
Assistant Research Professor, Dept. of Electrical & Computer Eng.  
Rutgers University  
Email: manish@cac.rutgers.edu

**Javier Diaz-Montes**, Ph.D.  
Assistant Professor, Dept. of Electrical & Computer Eng.  
Rutgers University  
Email: jdiaz@cac.rutgers.edu

**And many collaborators...**  
CometCloud: <http://cometcloud.org>

**THINKING PARALLEL: MULTI-CORES, VIRTUAL  
ELASTICITY, AND THE APPLICATION PROGRAMMER**  
**Geir Horn, University of Oslo, Norway**

CPUs can offer a large number of cores and dedicated accelerators. Large data centres offers large number of virtual machines in the Cloud. At the exception of some highly optimised eScience applications, these resources are used to support high throughput computing of data parallel applications. Mixing true parallelism with the cloud is inherently complicated, and the application programmer cannot think parallel. This talk will discuss the issue and present some possible ways forward, hopefully stimulating an interesting discussion on future research directions.

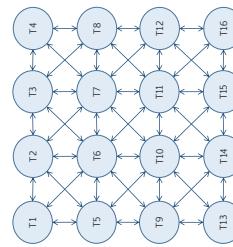


# Thinking parallel: Multi-cores, virtual elasticity, and the application programmer

Geir Horn



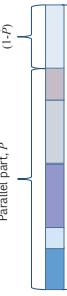
## The basics



Thinking parallel:  
Multi-cores, virtual elasticity,  
and the application programer



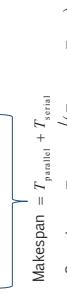
## Sequential scheduling (one processor)



## Makespan = $T_{\text{sequential}}$



## Makespan = $T_{\text{parallel}} + T_{\text{serial}}$



Task Interaction Graph (TIG)



## Speedup versus Scalability

Ideal speedup = number of processors =  $|N|$

Andahl's law<sup>3</sup>:

$$\text{Speedup} = \frac{1}{(1 - P) + \frac{P}{|N|}}$$

Max speedup as  $|N| \rightarrow \infty$  is

$$\frac{1}{1 - P}$$

Example: With  $P = 90\%$   
the max speedup is 10

Alternatively: Keep run time fixed, but increase problem size with the number of processors

Hypothetical sequential run time

$T_{\text{sequential}} = T_{\text{parallel}} + |N| \times T_{\text{parallel}}$

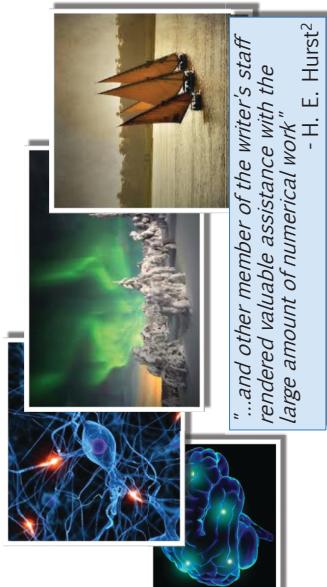
Gustafson-Barsis' law<sup>4</sup>: the scaled speedup is

$$\begin{aligned} SS &= T_{\text{sequential}} / (T_{\text{parallel}} + T_{\text{serial}}) \\ &= (T_{\text{serial}} + |N| \times T_{\text{parallel}}) / (T_{\text{parallel}} + T_{\text{serial}}) \\ &= |N| - \alpha(|N| - 1) \end{aligned}$$

with  $\alpha = T_{\text{serial}} / (T_{\text{parallel}} + T_{\text{serial}})$

[3] Gene M. Amdahl, "Validity of the single processor approach to solving large scale computing problems", *Journal of the ACM*, vol. 16, p. 308, Nov. 1968.

[4] Linda E. Gustafson and Kurt E. Barsis, "More effective use of parallel supercomputers", *Proceedings of the ACM SIGART Conference on Parallel Supercomputing*, May 1988.



"... and other member of the writer's staff rendered valuable assistance with the large amount of numerical work"  
- H. E. Hurst<sup>2</sup>

| 2

[1] Sir Alexander Hume, "The human brain as a steady state in numbers", *Front. Neurol.*, vol. 3, p. Article 31, Nov. 2012.  
[2] H. E. Hurst, "Long term storage capacity of reservoirs", *Transactions of the American Society of Civil Engineers*, pp. 770-799, 1951.

| 4

## Parallel digital computers



## More beyond Moore

Gordon E. Moore<sup>5</sup>:

*"...the number of transistors on a chip will double every 18 months"*

Mostly process scaling with 0.7 every 24 months from 0.8 µm in 1992 until ~2000

- Physical distance problems with scaling
- Voltage supply scaling and increased leaks
- Clock frequency: maximum around 5GHz
- Power density wall and heat
- Design complexity: exponential design team growth

[5] Gordon E. Moore, "Cramming More Components Onto Integrated Circuits," *Electronics*, pp. 115-117, Apr. 1965.

[6] J. M. Riedel, et al., "POWER system microarchitectures," *IBM Journal of Research and Development*, vol. 46, no. 3, pp. 502-518, Jun. 2002.

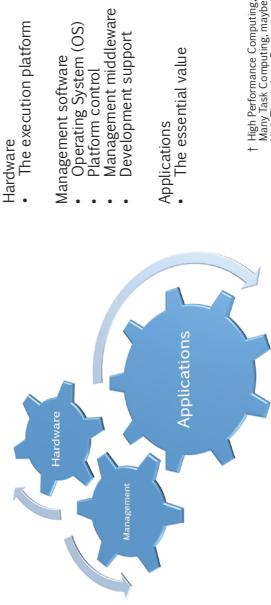
[7] J. M. Riedel, et al., "POWER system microarchitectures," *IBM Journal of Research and Development*, vol. 46, no. 3, pp. 502-518, Jun. 2002.

[8] L. A. Haas, et al., "8 Core 64 Thread Do-Double Execute SPARC-SOC: A Layer of Freedom," *Silicon Valley Computer Conference*, San Jose, CA, USA, Mar. 2009.

[9] L. A. Haas, et al., "8 Core 64 Thread Do-Double Execute SPARC-SOC: A Layer of Freedom," *Silicon Valley Computer Conference*, San Jose, CA, USA, Mar. 2009.

[10] J. M. Riedel, et al., "POWER System Microarchitectures," *IBM Journal of Research and Development*, vol. 46, no. 3, pp. 502-518, Jun. 2002.

## Parallel computing<sup>†</sup>



- High Performance Computing,  
Multi-core, Fast Compute,
- High Throughput Computing,  
not Dedicated Computing
- Development support

...on heterogeneous cores

Intel & AMD x86 dual-cores

High Performance Computing,  
Multi-core, Fast Compute

High Throughput Computing,  
not Dedicated Computing

Development support

Applications

The essential value

↑

"Assuming that this trend will follow Moore's Law scaling, mainstream systems will contain over 10 processing cores by the end of the decade, yielding unprecedented theoretical peak performance"

Justin Rattner  
Vice president and chief technology officer, Intel  
(2005)





Intel & AMD x86 dual-cores

High Performance Computing,  
Multi-core, Fast Compute

High Throughput Computing,  
not Dedicated Computing

Development support

Applications

The essential value

↑

"Assuming that this trend will follow Moore's Law scaling, mainstream systems will contain over 10 processing cores by the end of the decade, yielding unprecedented theoretical peak performance"

Justin Rattner  
Vice president and chief technology officer, Intel  
(2005)







↑

8 cores, 64 threads  
(2007)

↑

8 cores, 64 threads  
(2009)

↑

8 cores, 64 threads  
(2012)

[†] Multi-cores, virtual parallelism, and the application programmer interface.

[†] Multi-cores, virtual parallelism, and the application programmer interface.

[†] Multi-cores, virtual parallelism, and the application programmer interface.

[†] Multi-cores, virtual parallelism, and the application programmer interface.

[†] Multi-cores, virtual parallelism, and the application programmer interface.

[†] Multi-cores, virtual parallelism, and the application programmer interface.

[†] Multi-cores, virtual parallelism, and the application programmer interface.

[†] Multi-cores, virtual parallelism, and the application programmer interface.

[†] Multi-cores, virtual parallelism, and the application programmer interface.

[†] Multi-cores, virtual parallelism, and the application programmer interface.

[†] Multi-cores, virtual parallelism, and the application programmer interface.

## Hardware: The future

- > Dedicated accelerators<sup>15</sup> or proven FPGA templates or design patterns
  - > Proven library of modules with predictable performance<sup>15</sup>
    - Parameterised interconnects
    - Standard adapters and interfaces
    - Reusable building blocks
  - > Asymmetric speedup  $\geq$  symmetric speedup<sup>16</sup>
  - > "Software will take a more prominent role in the multi-core era"<sup>16</sup>
    - "...the programming model for heterogeneous architectures is much more complex"

## Operating Systems

- > "No current OS is 'really multithreaded'"<sup>18</sup>
  - > New approaches to operating systems
    - XtreamOS<sup>19</sup> = 
    - SLoSS = Service Oriented Operating Systems<sup>18</sup>
    - Tessilation<sup>19</sup>
  - > FUSE: OS support for easy HW accelerator integration<sup>20</sup>
  - > Native hypervisors + microkernels
  - > Real-time OS<sup>d</sup>
    - a Lutz Schubert, Universität Ulm  
b http://www.xtreamos.eu/ and http://research.cs.wisc.edu/cander/  
c http://www.slos-project.eu/  
d QNX Neutrino RTOShttp://www.qnx.com/products/neutrinos/rta.html or  
INTEGRITY http://www.qnx.com/products/rtos/integrity.html

<sup>15</sup> John Stark, "The new landscape of parallel computer architecture," *Journal of Physics: Conference Series*, vol. 78, p. 012004, Jul. 2007.

<sup>16</sup> David R. Kaeli, "Multi-core design automation challenges - a view of the state-of-the-art," *Proceedings of the 4th annual Design Automation Conference (DAC '07)*, San Diego, California, USA, 2007.

<sup>17</sup> Steve Ake, "A Survey of the State-of-the-Art in Multi-core Processor Design," *IEEE Transactions on Parallel and Distributed Systems*, vol. 19, no. 12, pp. 1725-1737, Dec. 2008.

<sup>18</sup> Steve Ake, "A Survey of the State-of-the-Art in Multi-core Processor Design," *IEEE Transactions on Parallel and Distributed Systems*, vol. 19, no. 12, pp. 1725-1737, Dec. 2008.

<sup>19</sup> Steve Ake, "A Survey of the State-of-the-Art in Multi-core Processor Design," *IEEE Transactions on Parallel and Distributed Systems*, vol. 19, no. 12, pp. 1725-1737, Dec. 2008.

<sup>20</sup> Steve Ake, "A Survey of the State-of-the-Art in Multi-core Processor Design," *IEEE Transactions on Parallel and Distributed Systems*, vol. 19, no. 12, pp. 1725-1737, Dec. 2008.

## High Throughput Mass Market

- > Desktop
  - Legacy software "does the job"
  - Current software developed for single-core
  - Per-application performance is important (only) if the load consists of only a few applications or if there are performance-critical applications<sup>33</sup>
- > Servers
  - Database and web servers are designed for high throughput
  - Idle time can be masked by multi-threading<sup>33</sup>
  - > Large parallel application part = more cores, otherwise more complex cores<sup>16</sup>

<sup>33</sup> Pin-limits: no technology in sight!<sup>33</sup>

> Three scenarios<sup>35</sup>

- "Drop the ball"  $\Rightarrow$  Cloud computing
- "Niche markets"  $\Rightarrow$  Multimedia, gaming...
- "Scalable, dependable, software development"<sup>13</sup>

> Urgently needed: **"Parallel computing for all"**<sup>TM</sup>

- Standardised, industry accepted development platforms
- Education of programmers on these platforms
- Requirements engineering for parallelism

<sup>35</sup> [33] David Pearson, "The trend with multi-core," *IEEE Spectrum*, vol. 47, no. 7, pp. 28-32, 25-54, Jul. 2010.

<sup>15</sup> [34] David Pearson, "The trend with multi-core," *IEEE Spectrum*, vol. 47, no. 7, pp. 28-32, 25-54, Jul. 2010.

<sup>12</sup>

## Conclusion

- > Thinking parallel, virtual elasticity, and the application programmer interface

<sup>13</sup>

<sup>14</sup>

<sup>15</sup>

<sup>16</sup>

<sup>17</sup>

<sup>18</sup>

<sup>19</sup>

<sup>20</sup>

<sup>21</sup>

<sup>22</sup>

<sup>23</sup>

<sup>24</sup>

<sup>25</sup>

<sup>26</sup>

<sup>27</sup>

<sup>28</sup>

<sup>29</sup>

<sup>30</sup>

<sup>31</sup>

<sup>32</sup>

<sup>33</sup>

<sup>34</sup>

<sup>35</sup>



*In a soldier's stance, I earned my hand  
At the mangled dogs who teach  
Fearing not that I'd become my enemy  
In the instant that I breach*

*My pathway fed by conviction beats*

*Many from stern to bow*

*Ah, but I was so much older then*

*Im younger than that now*

*Yes, my gaunted stock had when abstract threats*

*To us, To us,*

*Deceived me into thinking*

*I had something to protect*

*Good and bad, define these terms*

*Quite clear, no doubt, somehow*

*Ah, but I was so much older then*

*Im younger than that now*

**Bob Dylan**

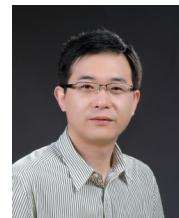
*My back pages*

Geir Horn  
[Geir.Horn@mnn.uio.no](mailto:Geir.Horn@mnn.uio.no)  
+47 93 05 93 35



**SIMPLIFIED CLOUD CONTROL USING DIMENSION  
REDUCTION****Jianguo Yao, Shanghai Jiao Tong University**

Automated management of complex information technology applications such as cloud systems requires dynamic configuration of both application-level and system-level parameters. The existence of large number of tunable parameters makes it difficult to design a feedback controller that adjusts these parameters effectively in order to achieve the application-level performance targets. In this talk, we will summarize our recent work which introduces a new approach for simplified control architecture of large-scale complex systems based on dimension reduction techniques. It combines the online selection of critical control knobs through LASSO -- a powerful L1 -constrained fitting method, and Compressive Sensing (CS)-- a L1-optimization method, and the design of adaptive control of the identified knobs. We use evaluation results to demonstrate the effectiveness of this new approach.





## Outline

- > Motivation
- > Problems and Challenges
- > Solution architecture
- > Dimension Reduction Methods
  - ⦿ LASSTOLARS
  - ⦿ Compressive Sensing
- > Three-module controller design
  - ⦿ Dimension reduction
  - ⦿ RLS-based model identification
  - ⦿ Linear quadratic optimal controller
- > Evaluation results
  - ⦿ Simulation
  - ⦿ Experiment
  - ⦿ Conclusion and future work

2/21

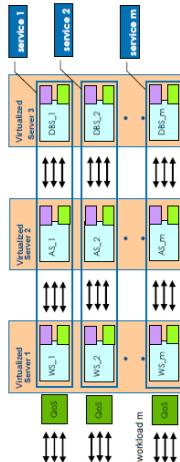
# Simplified Cloud Control Using Dimension Reduction

Jianguo Yao (SJTU), Xue Liu(McGill University),  
Joint work with Xiaoyun Zhu (VMware)

## Motivation

Increasingly complex cloud systems

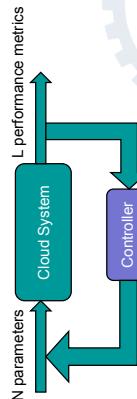
- > Cloud system and applications are getting more complex
- > Large numbers of system-level and application-level parameters
  - ⦿ System level: no. of VMs, resource allocation (CPU, GPU, memory, network, I/O), workload/demand placement, cache size, processor frequency
  - ⦿ Application-level: no. of processes/thread, no. of database connections, time-out, keepalive



3/21

## Problems

- > Problem
  - ⦿ How to adjust these parameters to achieve application-level SLOs?
  - ⦿ Manual tuning is hard  $\rightarrow$  automation via feedback control



4/21

**Cloud Control** needs a **model** for cloud system  
to design the controller.

## Challenges

- Challenges
    - Too many tunable parameters (inputs)  $\rightarrow$  which ones are the most critical?
    - Set of critical parameters may vary over time
    - A controller that tunes all the parameters is not efficient due to the high dimension
- $N$  parameters  $\xrightarrow{L}$  performance metrics
- $M < N$
- $M$  Selected parameters  $\xrightarrow{L}$  performance metrics
- How to automatically build a simplified model with low dimension.

5/21

## A motivating example A realistic case

- To adjust response time in Apache, we can insert a function of Sleep call.
- However, there are many functions in the Apache in which only a few can dramatically effect the response time.

TABLE 1  
Coefficients for functions.

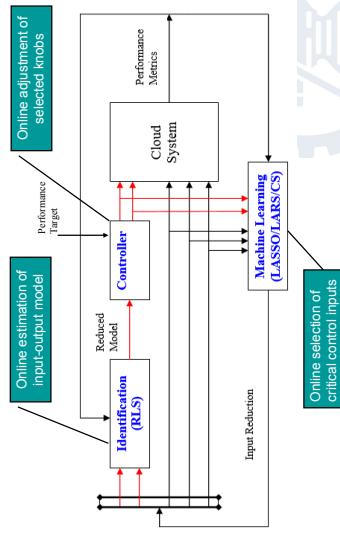
Coefficients	More than 100	Functions	Coefficients
$r(k+1) = \sum_{i=1}^m \lambda_i s_i(k) + e_1(k)$			$\lambda_1 = 33.6399$
			$\lambda_2 = -0.00123$
			$\lambda_3 = -0.001$
			$\lambda_4 = -0.0008$
			$\lambda_5 = 0.0006$
			$\lambda_6 = 0.1180$
			$\lambda_7 = -0.0323$
			$\lambda_8 = 0.0164$
			$\lambda_9 = -0.0806$
			$\lambda_{10} = -0.2022$

$$r(k+1) = \lambda_1 s_1(k) + e_1(k)$$

$$\text{where } e(k) = \sum_{i=2}^m \lambda_i s_i(k) + e_1(k)$$

6/21

## Simplified cloud control architecture



7/21

## Dimension reduction using Lasso

- Lasso – Shrinkage and selection method for linear regression

☞ R. Tibshirani, "Regression Shrinkage and Selection via the Lasso," *Journal of Royal Statistics Society, 1996*.  
<http://statweb.stanford.edu/~tibs/lasso.html>

System inputs:  $u_1, u_2, \dots, u_m$ , zero mean and unit length

System output:  $y$ , zero mean

Find model:  $\hat{y} = \sum_{j=1}^m \hat{x}_j u_j$

$$\text{where } \min_{\hat{y}} \|\hat{x}\|_1 = \sum_{j=1}^m |\hat{x}_j| < t$$

Tuning parameter

8/21

## Dimension reduction using CS

- Measurement data  $y = \Phi f$ ,  $M < N$
- $y$    $f$  
- $\Phi$  
- $y$    $f$  
- $\Phi$  
- Applying Compressive Sensing (CS) when : signal  $x$  is compressible, sparse...
- $y$    $x$  
- $\Phi$  
- $y$    $x$  
- $\Phi$  
- Construct

## Recovery

- Recovery
- $y = \Phi f$ ,  $f = \Psi x$
- $y$    $f$  
- $\Phi$  
- $\Psi$  
- Solve for  $x$ , s.t.  $y = \Phi \Psi x$
- Optimization problem of CS
- $\min_x \|x\|_1$
- $M \ll N$
- $\text{min}_x \|x\|_1$
- s.t.  $\|y - \Phi \Psi x\|_2 \leq \varepsilon$
- Our problem
- $\hat{y} = \sum_{j=1}^m \hat{x}_j u_j$

10/21

9/21

## An Example Estimation of coefficients

- Least Angle Regression (LARS) method allows selection of  $s < m$  inputs to predict the output

- ⦿ A variant of LASSO, a simpler method for computation
- ⦿ B. Efron, I. Johnstone, T. Hastie, and R. Tibshirani, "Least angle regression," *Annals of Statistics*, 2004.

### Input-output model

$$\begin{aligned} y(k+1) = & 1.0 u_1(k) + 0.1 u_2(k) + 0.1 u_3(k) \\ & + 2.0 u_4(k) + 0.1 u_5(k) + 0.1 u_6(k) \\ & + 5.0 u_7(k) + e(k) \end{aligned}$$

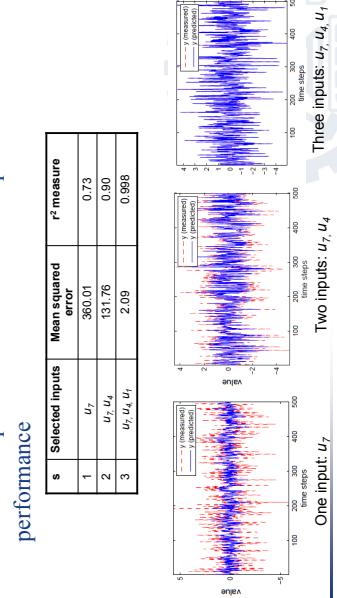
### Estimation using Compressive Sensing

$$\hat{x} = [0.9711 \ 0.0066 \ -0.0013 \ 2.0068 \ 0.0250 \ 0.0007 \ 4.9945]^T$$

11/21

## An Example Estimation of coefficients

- Selected input subsets V.S.  $n$  and their prediction performance



12/21

Three inputs:  $u_7, u_4, u_1$   
Two inputs:  $u_7, u_4$   
One input:  $u_7$

## Online model estimation

- Input-output model using selected inputs

$u_s$  : vector of  $s$  selected inputs

$$y(k+1) = X\phi(k) + e(k)$$

$$\phi(k) = u_s(k).$$

$$\begin{aligned} \hat{X}(k+1) &= \hat{X}(k) + \frac{e(k+1)\phi^T(k)P(k-1)}{\lambda + \phi^T(k)P(k-1)\phi(k)}, \\ P^{-1}(k) &= P^{-1}(k-1) \left( 1 + (\lambda - 1) \frac{\phi^T(k)P(k-1)\phi(k)}{[\phi^T(k)\phi(k)]^2} \right) \phi(k)\phi^T(k), \end{aligned}$$

- Online adaptation using recursive least squares (RLS) with exponential forgetting  $\lambda$
- Estimation of  $X$  updated in every interval  $k$

13/21



## Linear quadratic optimal controller

- Based on the estimated model with reduced dimension

$$y(k+1) = X\phi(k) + e(k)$$

- Minimizing quadratic cost function

$$J = \|W(y(k+1) - y_{ref}(k+1))\|^2 + \|Q(u_s(k) - u_s(k-1))\|^2$$

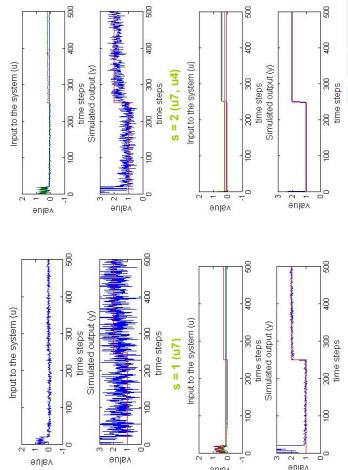
- Optimal solution:

$$u_s^*(k) = ((W\hat{X}(k))^T W\hat{X}(k) + Q^T Q)^{-1}(W\hat{X}(k)^T W y_{ref}(k+1) + Q^T Qu_s(k-1))$$

14/21



## Simulation Varying number of selected inputs



15/21

## Simulation varying the system behavior

- At interval k = 250, system model changes to

$$y(k+1) = 0.1 u_1(k) + 3.1 u_2(k) + 14.1 u_3(k) + 0.1 u_4(k)$$

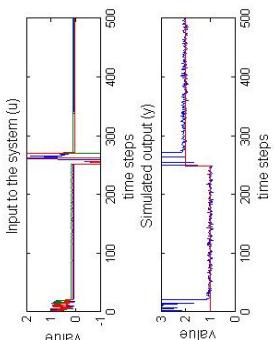
$$+ 2.1 u_5(k) + 0.1 u_6(k) + 0.1 u_7(k) + e(k)$$

- Controller detects the degradation of performance, and starts to collect 20 new samples of input-output data

- At interval k = 270, LARS selects a new set of inputs {u<sub>2</sub>, u<sub>3</sub>, u<sub>5</sub>}

16/21

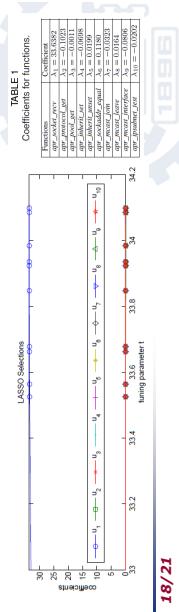
## Simulation varying the system behavior



17/21

## Experiment A realistic case

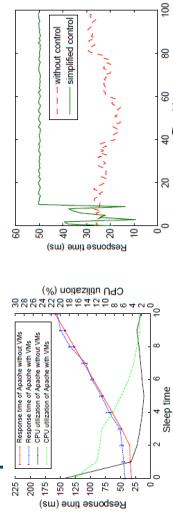
- A tested with iCore7-2600k 3.4GHz CPU and 16GB RAM.
- Each hosted VM is assigned dual cores and 2GB RAM. We choose Windows<sup>7</sup> x64 operating system for both host OS and guest OSes.
- Simplified Model:  $r(k+1) = \lambda_1 s_1(k) + e(k)$
- We vary the tuning parameter "r" from 33 to 34.2
- Estimates of regression coefficients using LRAS/LASSO for the realistic case



18/21

## Experiment A realistic case

- Estimated model is piece-wise linear.
- Set 50.0 ms as the reference.
- The average response time of Apache has been controlled at 49.943 ms.
- Response time of Apache with/without reduced dimension simplified control.



19/21

## Related work

- Prior work on applying control theory to computing systems
  - ② Control knobs determined in advance
- Xiong *et al.*, ICDCS'11, ICPE'2013
  - ② Automated model-driven framework for application performance diagnosis in consolidated Cloud Environments
  - ② Robust provisioning of N-Tier cloud workloads: a multi-level control approach
- Diao *et al.*, IM'03
  - ② Online discovery of critical metrics for a database system, used to construct a quantitative model
  - ② No online adaptation of the model or the metrics
- Classical model reduction in control theory
  - ② Reduces the dimensionality of the state space

20/21

## Conclusions and future work

- First framework for combined online selection of control knobs and dynamic tuning of the knobs
- Three-module controller achieves three design objectives
  - ② Online selects a subset of control knobs that have the most significant impact on the controlled output
  - ② Dynamically tunes the selected control knobs effectively to maintain the system output at the desired value
  - ② Automatically detects the change in the most critical knobs and uses the new knobs to regulate the system output accordingly
- Next step
  - ② Apply the framework to a real problem in large-scale cloud systems management
    - Dynamic allocation of multiple resources
    - Application configuration management

21/21

## Reference

- 
- Jianguo Yao, Xue Liu, Xiaoyun Zhu and Haibing Guan, "Control of Large-Scale Systems Through Dimension Reduction", IEEE Transactions on Service Computing, 2014. (Preprint)
  - Jianguo Yao, Xue Liu, Xiaoyun Zhu, "Reduced Dimension Control Based on Online Recursive Principal Component Analysis", in Proceedings of the 2009 American Control Conference (ACC'2009), St. Louis, MO, USA, 2009.
  - T. Abdelzher, J.A. Stankovic, C. Lu, R. Zhang, and Y. Lu, "Feedback performance control is software services", IEEE Control System Magazine, vol. 23, no. 3, pp.74-90, 2003.
  - R. Tibshirani, "Regression Shrinkage and Selection via the Lasso," Journal of Royal Statistics Society, 1996.
  - R.G.Baraniuk, "Compressive Sensing [Lecture Notes]," IEEE Signal Processing Magazine, vol.24, no.4, pp.118-121, 2007.
- 

Questions?

Thanks!

## Backup



## Architecture

```
> Goal: Ensure FPSVW = Target  
> Insert sleep function       $T_{sleep} = \frac{1000}{F} - T_{cpu} - T_{gpu}$   
  {  
    DrawShapes(&VGA_Buffer);  
    SleepUntil(1.1f); GPU to  
    display the buffered content.  
  }
```



Challenge: Uncertainties from execution time of CPU and GPU

Solution: Using feedback control to address the uncertainties

**REAL-TIME PERFORMANCE CONTROL OF ELASTIC****VIRTUALIZED NETWORK FUNCTIONS****Tommaso Cucinotta, Bell Labs, Alcatel-Lucent, Ireland**

Controlling end-to-end service quality, and particularly real-time performance and reliability, of time-sensitive applications in cloud environments is overly challenging. The problem is even harder when trying to switch from the earlier world of network functions shipped as physical boxes, to the emerging paradigm of virtualized cloud-ready elastic network functions, where still the "5 9s" and tight sub-second real-time performance requirements as coming from precise SLAs have to be met. This talk provides an overview of past and ongoing research carried out at Bell Labs on these challenging issues.



# Real-time Performance Control of Elastic Virtualized Network Functions

**Tomaso Cucinotta**  
Bell Laboratories, Alcatel-Lucent  
Dublin, Ireland

## Introduction

### A new era of computing for ICT

- Wide availability of broadband connections
- ==> shift in computing paradigms towards distributed computing (**cloud computing**)
  - More and more resources provided remotely
  - Not only *remote storage and batch processing*
  - But also *remote processing for interactive applications*
  - Network operators are shifting provisioning of critical network services to virtualized network functions (through **private or hybrid cloud** provisioning models)

### Examples

- **Virtual Reality** with heavyweight physics simulations
- Distributed editing of HD video, **film post-production**, **Alcatel-Lucent**

Tommaso Cucinotta - Bell Laboratories - Dublin

Copyright © 2012 ALACATEL-LUCENT. ALL RIGHTS RESERVED

AT THE SPEED OF IDEAS™

Tommaso Cucinotta - Bell Laboratories - Dublin

Tommaso Cucinotta - Bell Laboratories - Dublin

Copyright © 2012 ALACATEL-LUCENT. ALL RIGHTS RESERVED

AT THE SPEED OF IDEAS™

Tommaso Cucinotta - Bell Laboratories - Dublin

## Introduction

**Alcatel-Lucent**   
Tommaso Cucinotta - Bell Laboratories - Dublin

## Introduction

Virtualization technologies are key

- For **IaaS** providers (Cloud Computing)
- For **server consolidation**
- Different virtualization technologies



...



Tommaso Cucinotta - Bell Laboratories - Dublin

Tommaso Cucinotta - Bell Laboratories - Dublin

Copyright © 2012 ALACATEL-LUCENT. ALL RIGHTS RESERVED

AT THE SPEED OF IDEAS™

Tommaso Cucinotta - Bell Laboratories - Dublin

Tommaso Cucinotta - Bell Laboratories - Dublin

Copyright © 2012 ALACATEL-LUCENT. ALL RIGHTS RESERVED

AT THE SPEED OF IDEAS™

Tommaso Cucinotta - Bell Laboratories - Dublin

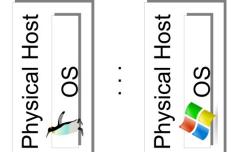
## Introduction

Virtualization technologies are key

- For **IaaS** providers (Cloud Computing)

- For **server consolidation**

### Different virtualization technologies



AT THE SPEED OF IaaS™

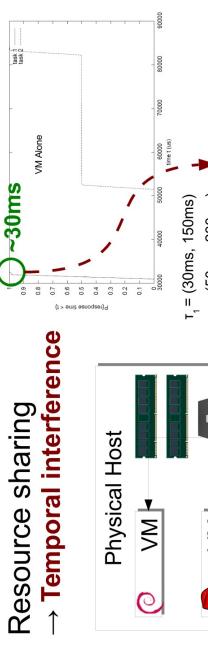
COPYRIGHT © 2012 ALCATEL-LUCENT. ALL RIGHTS RESERVED

Tommaso Cicalini - Bell Laboratories - Dublin

Alcatel-Lucent

## Need for Performance Isolation

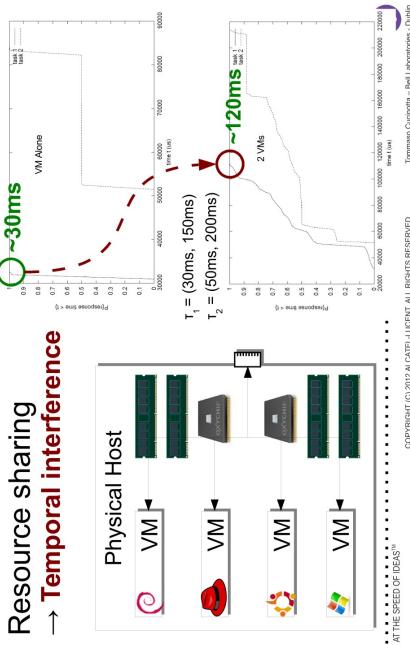
Resource sharing  
→ Temporal interference



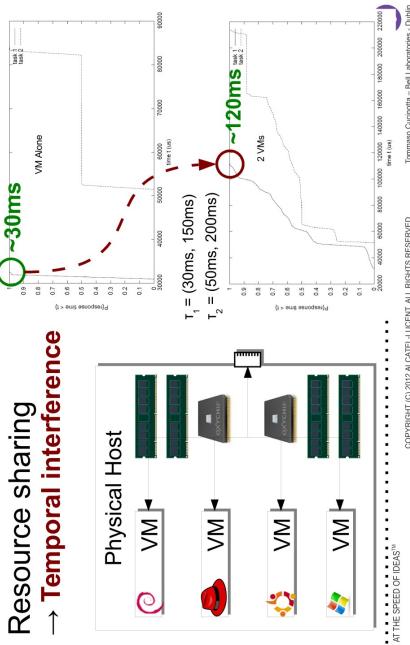
PARTICIPANTS

81

Resource sharing  
→ Temporal interference



Resource sharing  
→ Temporal interference



AT THE SPEED OF IaaS™

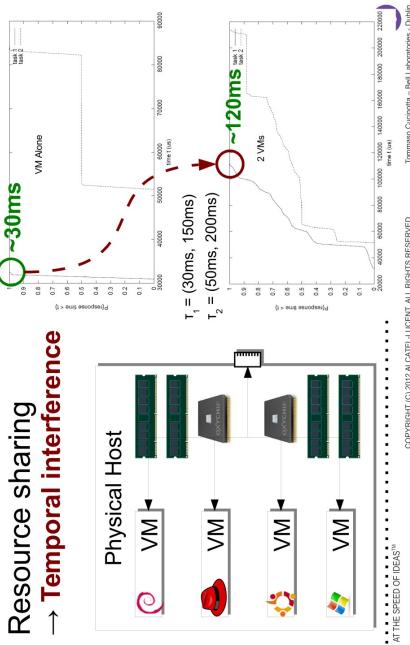
COPYRIGHT © 2012 ALCATEL-LUCENT. ALL RIGHTS RESERVED

Tommaso Cicalini - Bell Laboratories - Dublin

Alcatel-Lucent

## Need for Performance Isolation

Resource sharing  
→ Temporal interference



AT THE SPEED OF IaaS™

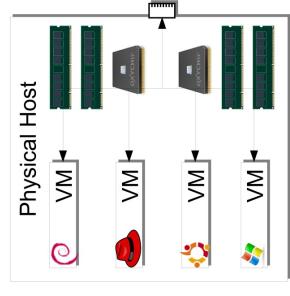
COPYRIGHT © 2012 ALCATEL-LUCENT. ALL RIGHTS RESERVED

Tommaso Cicalini - Bell Laboratories - Dublin

Alcatel-Lucent

## Need for Performance Isolation

Resource sharing  
→ Temporal interference



AT THE SPEED OF IaaS™

COPYRIGHT © 2012 ALCATEL-LUCENT. ALL RIGHTS RESERVED

Tommaso Cicalini - Bell Laboratories - Dublin

Alcatel-Lucent

## Co-Scheduling Virtual Machines

### Issues in deploying RT SW in VMs

- Scheduling and timing
  - VM scheduling impacts on the vision of time by guest OSes**
    - Time granularity (for measuring time and setting timers)
    - Non-uniform progress-rate of applications
    - SMP-enabled guests
      - Spin-lock primitives assume release of locks within very short time-frames
    - What happens if the **lock-owner VM is descheduled**?

### Benchmarking

- A VM may be deployed on different HW (SOA scenario)
  - How to achieve predictable performance?
- VMs may be deployed on **General-Purpose HW** (with cache)
  - How to account for **HW-level interferences**?

## Possible Solutions

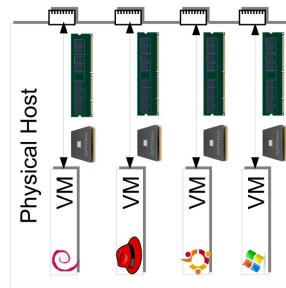
Hardware replication and  
static partitioning

- Computing
  - Multi-core (1 core per VM)**
- Networking
  - Multiple network adapters (1 network adapter per VM)**
  - Multi-queue adapters

- Drawbacks
  - Limitation of **flexibility**
  - Under-utilization** of resources



AT THE SPEED OF IDEAS™  
Copyright © 2012 ALCATEL-LUCENT. ALL RIGHTS RESERVED  
Tarmanno Giurato - Bell Laboratories - Dublin



Another approach

- Let **multiple VMs use the same resources**
  - Proper management of **priority inversion**
  - Reduced overheads (limited number of preemptions)
  - Run-time schedulability analysis and **admission control**
- For example
  - Computing
    - Xen credit-based, SEDF schedulers, RT-Xen exts
  - Networking
    - QoS-aware protocols (IntServ, MPLS)
- Advantages
  - Increased **flexibility**
  - Increased **resource saturation levels**
  - Reduced infrastructure costs**



## Possible Solutions

- Proper management of **shared resources**: what MP **resource-sharing protocol** is appropriate?
  - Proper management of **priority inversion**
  - Reduced overheads (limited number of preemptions)
  - Run-time schedulability analysis and **admission control**
- Proper management of **shared resources**: what MP **resource-sharing protocol** is appropriate?
  - Proper management of **priority inversion**
  - Reduced overheads (limited number of preemptions)
  - Run-time schedulability analysis and **admission control**

## Co-Scheduling Virtual Machines

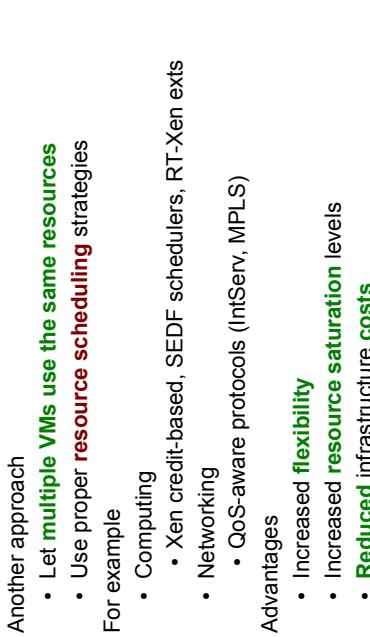
### Issues in deploying RT SW Components in VMs

- Temporal isolation** across VMs
  - Compute-bound and I/O-bound VMs
  - Shared host resources (e.g., network interrupt drivers)
  - Intensive I/O on virtualised peripherals (big-data)
- Proper management of **shared resources**: what MP **resource-sharing protocol** is appropriate?
  - Proper management of **priority inversion**
  - Reduced overheads (limited number of preemptions)
  - Run-time schedulability analysis and **admission control**



AT THE SPEED OF IDEAS™  
Copyright © 2012 ALCATEL-LUCENT. ALL RIGHTS RESERVED  
Tarmanno Giurato - Bell Laboratories - Dublin

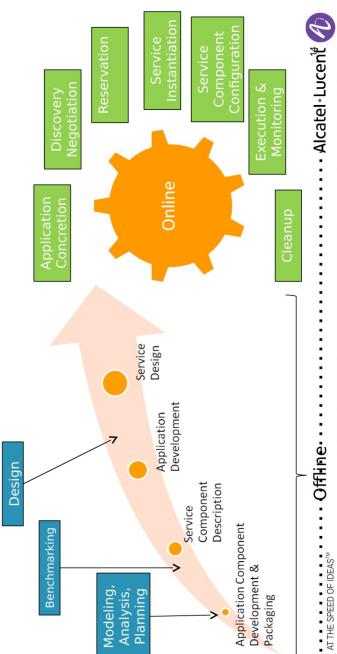
- Proper management of **shared resources**: what MP **resource-sharing protocol** is appropriate?
  - Proper management of **priority inversion**
  - Reduced overheads (limited number of preemptions)
  - Run-time schedulability analysis and **admission control**
- Proper management of **shared resources**: what MP **resource-sharing protocol** is appropriate?
  - Proper management of **priority inversion**
  - Reduced overheads (limited number of preemptions)
  - Run-time schedulability analysis and **admission control**



AT THE SPEED OF IDEAS™  
Copyright © 2012 ALCATEL-LUCENT. ALL RIGHTS RESERVED  
Tarmanno Giurato - Bell Laboratories - Dublin

## IRMOS Two-Phase Approach

### General IRMOS Approach



### Approach

Traditional (hard) real-time techniques are not appropriate

- lead to poor resource utilization
- imply high/unsustainable development costs

Soft real-time techniques are more appropriate

• Stochastic models for system/QoS evolution

• Probabilistic guarantees (as opposed to deterministic ones)

Pragmatic approach

- Theory is always applied
- on real **GPOS** (Linux)
  - with a real **Virtual Machine Monitor** (KVM)
  - on real **multimedia applications** (mp3player, vlc, ...)

### Basic Building blocks

- Linux / KVM enriched with our RT Scheduler(s)
- Each VMU is attached RT scheduling parameters (defining its temporal capsule)
- Improvements on the real-time virtualization performance
  - Modifications at the hypervisor level
  - Modifications at the kernel level

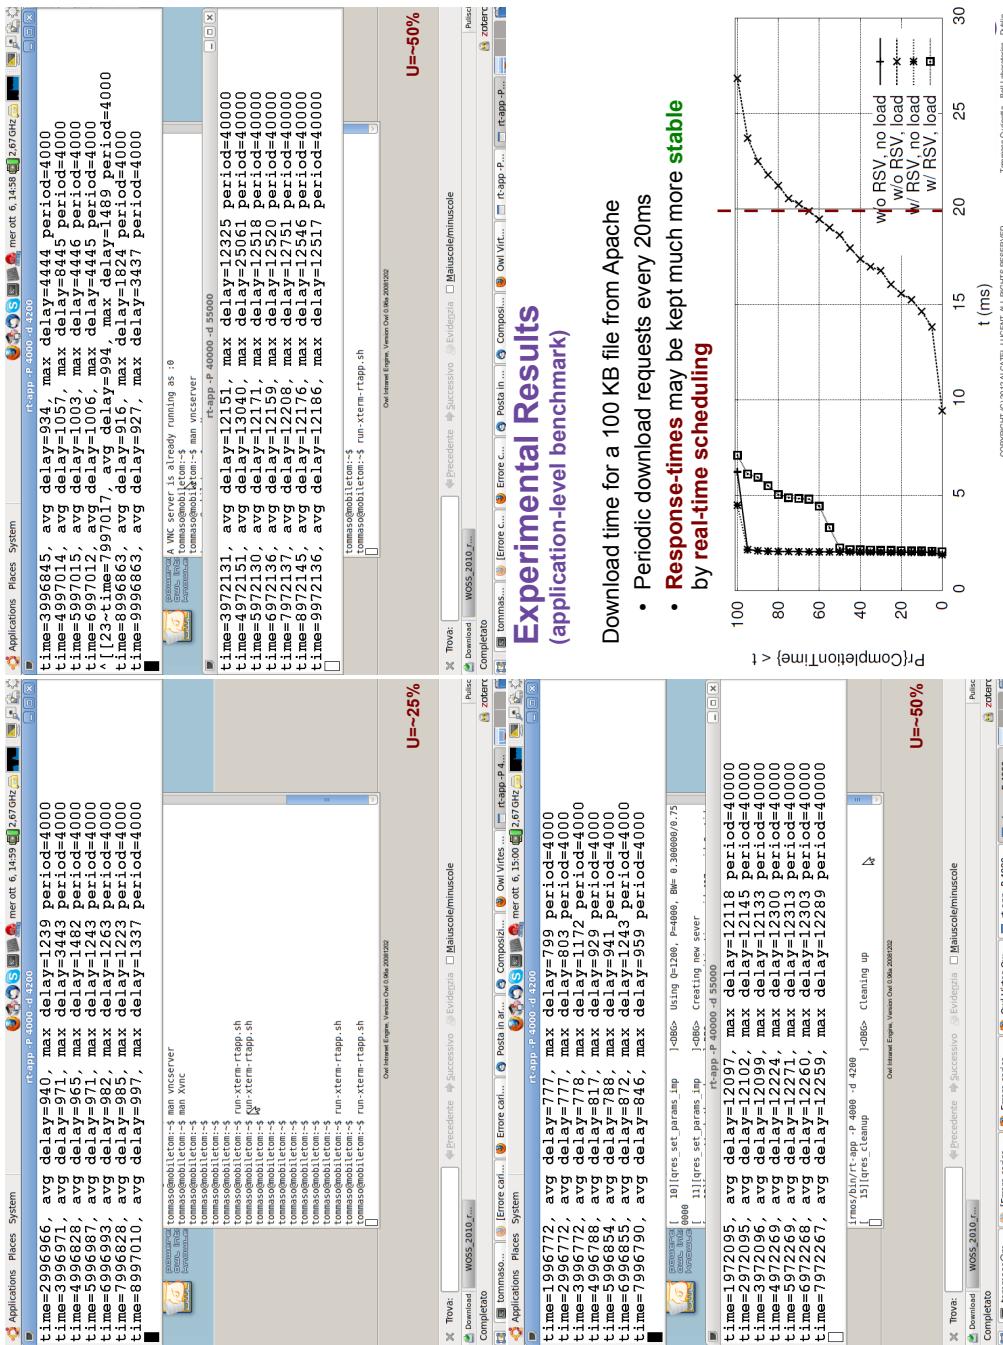
• Analysis of Virtualized RT applications by Hierarchical Real-Time Schedulability Analysis



16

Tommaso Giordano - Bell Laboratories - Dakar

Tommaso Giordano - Bell Laboratories - Dakar



## Plethora of Cloud Providers, Tools and Frameworks

- Cloud IaaS
  - Amazon, Rackspace, Google Compute, ...
  - OpenNebula, OpenStack, CloudStack
  - CloudBand, ...
- Configuration Management (skip)
- **Monitoring and Orchestration**
  - Amazon AutoScaling, Heat+CloudFormation, Cloudify, CloudFoundry, Chef Recipes, ...

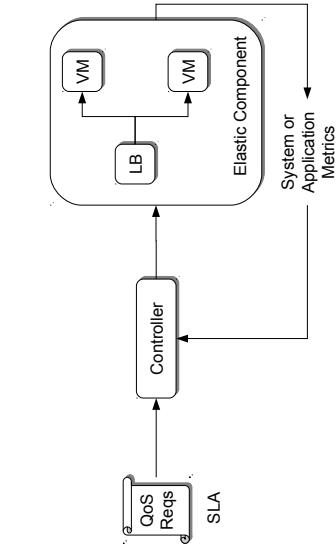
## Controlling Elastic Virtualized Applications



At the speed of IaaS™  
Copyright © 2012 ALCATEL-LUCENT. All rights reserved.  
Tommaso Crociato - Bell Laboratories, Darmstadt

## Elasticity Loop

But...



Adaptation logic built  
on unstable terrain!



At the speed of IaaS™  
Copyright © 2012 ALCATEL-LUCENT. All rights reserved.  
Tommaso Crociato - Bell Laboratories, Darmstadt



At the speed of IaaS™  
Copyright © 2012 ALCATEL-LUCENT. All rights reserved.  
Tommaso Crociato - Bell Laboratories, Darmstadt



At the speed of IaaS™  
Copyright © 2012 ALCATEL-LUCENT. All rights reserved.  
Tommaso Crociato - Bell Laboratories, Darmstadt

At the speed of IaaS™  
Copyright © 2012 ALCATEL-LUCENT. All rights reserved.  
Tommaso Crociato - Bell Laboratories, Darmstadt

## But...

Adaptation logic built  
on unstable terrain!



Can we make  
anything better?



AT THE SPEED OF IDEAS™  
Copyright © 2012 ALCATEL-LUCENT. ALL RIGHTS RESERVED

• • • • • Alcatel-Lucent

Termao Gómez - Bell Laboratories - Dublin

AT THE SPEED OF IDEAS™  
Copyright © 2012 ALCATEL-LUCENT. ALL RIGHTS RESERVED

• • • • • Alcatel-Lucent

Termao Gómez - Bell Laboratories - Dublin

## Related Publications

- "Elastic Admission Control for Federated Cloud Services," (to appear on) IEEE Transactions on Cloud Computing
- "Data Centre Optimisation Enhanced by Software Defined Networking," (to appear) in IEEE CLOUD 2014
- "Brokering SLAs for end-to-end QoS in Cloud Computing," CLOSER 2014, Barcelona
- "End-to-End Service Quality for Cloud Applications," GECON 2013, Zaragoza
- "Run-time Support for Real-Time Multimedia in the Cloud," REACTION 2013, Vancouver
- "Admission Control for Elastic Cloud Services," IEEE CLOUD 2012, Hawaii
- "Virtualised e-Learning with Real-Time Guarantees on the iRMOS Platform," IEEE SOCA, December 2010 [best paper award]
- "Hierarchical Multiprocessor CPU Reservations for the Linux Kernel," OSPERT 2009, Dublin

## Thanks for your attention

Questions ?

AT THE SPEED OF IDEAS™  
Copyright © 2012 ALCATEL-LUCENT. ALL RIGHTS RESERVED

• • • • • Alcatel-Lucent

Termao Gómez - Bell Laboratories - Dublin

## AN ADAPTIVE UTILISATION ACCELERATOR FOR VIRTUALIZED ENVIRONMENTS

Giovanni Toffetti, IBM Haifa Research Lab

One of the key enablers of a cloud provider competitiveness is ability to over-commit shared infrastructure at ratios that are higher than those of other competitors, without compromising non-functional requirements, such as performance. A widely recognized impediment to achieving this goal is so called "Virtual Machines sprawl", a phenomenon referring to the situation when customers order Virtual Machines (VM) on the cloud, use them extensively and then leave them inactive for prolonged periods of time. Since a typical cloud provisioning system treats new VM provision requests according to the nominal virtual hardware specification, an often occurring situation is that the nominal resources of a cloud/pool become exhausted fast while the physical hosts utilization remains low.

We present IBM adaPtive UtiliSation AcceleratoR (IBM PUL-SAR), a cloud resources scheduler that extends OpenStack Nova Filter Scheduler. IBM PULSAR recognises that effective safely attainable over-commit ratio varies with time due to workloads' variability and dynamically adapts the effective over-commit ratio to these changes.



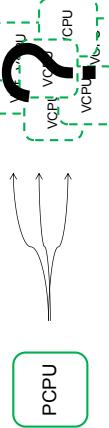
## An Adaptive Utilization Accelerator for Virtualized Environments

### LCCC Workshop in Cloud Control

David Breitgand, Zvi Dubitzky, Amir Epstein, Oshnit Feder, Alex Glikson, Inbar Shapira and Giovanni Toffetti  
Cloud Operating Systems Technologies – IBM Haifa Research Lab



### Common solution: resource over-commit



$$\text{Over-commit Ratio (OCR)} = \# \text{VCpus} / \# \text{PCpus}$$

- OCR is a cloud-wide configuration parameter
- Default CPU OCR for Openstack: 16()

OCR => infrastructure utilization  
 OCR => risk of congestion

performance degradation

### What is the problem? VM sprawl in private clouds

- VM sprawl
  - Proliferation of inactive / unused VMs in clouds
  - Stems from cloud provisioning model (and relative lack of control)
- In the absence of resource utilization models, VM provisioning is based on nominal resource demand

Looking at a VM:

2 vCPUs  
 4 GB RAM

Nominal resource demand

Looking at a DC:

100% SPAM!

My private cloud

Fully utilized nominal capacity

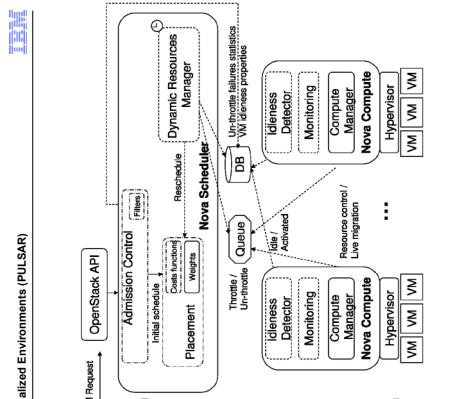
VM provisioning based on nominal resource demand  
 Unsatisfied VM placement demand

### Our proposal – Adaptive over-commit

- RATIONALE:
  - There is **NO “right” fixed OCR**
  - VMs “activity” vs. “idleness” are application-dependent and vary over time
  - Need for automated solution

#### ▪ GOALS/FEATURES:

- Increase DC utilization
- Minimize performance degradation
- Transparent to VM tenants
- No assumptions/forecast on VM resource consumption



## Pulsar Implementation

- Full implementation as extension to OpenStack Havana release (soon IceHouse)
- Idleless detector running on each host
- Adjusted capacity filter
- Dynamic resource manager
- Admission control mechanism



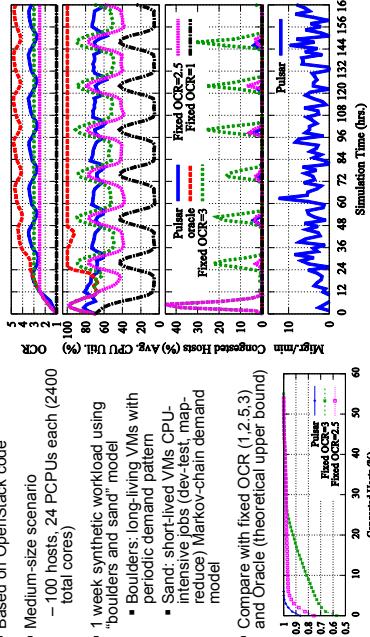
**Openstack**  
CLOUD SOFTWARE

© 2014 IBM Corporation

## IBM Adaptive Utilization Accelerator for Virtualized Environments (PULSAR)

### Synthetic workload simulation

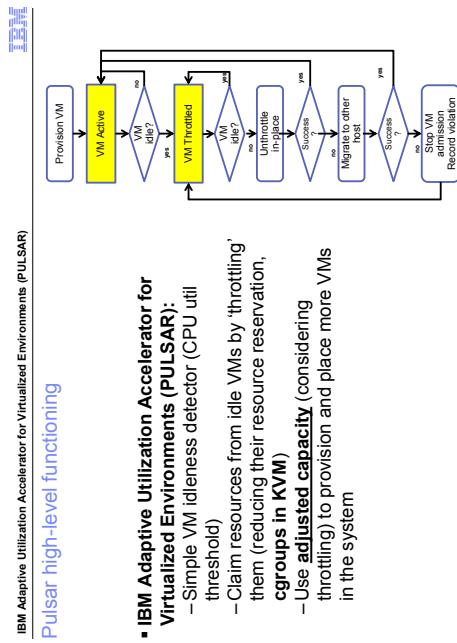
- Based on OpenStack code
- Medium-size scenario
  - 100 hosts, 24 PCPUs each (2400 total cores)
- 1 week synthetic workload using "boulders and sand" model
  - Boulders: long-living VMs with periodic demand pattern
  - Sand: short-lived VMs CPU-intensive jobs dev-test, map-reduce) Markov-chain demand model
- Compare with fixed OCR (1.2.5.3), and Oracle (theoretical upper bound)



## IBM Adaptive Utilization Accelerator for Virtualized Environments (PULSAR)

### Pulsar evaluation

- Experiments
  - Smaller testbed
  - Full implementation
  - (Synthetic) trace-driven workload: boulders and sand model
  - Measure performance degradation
- Simulations
  - Large testbeds
  - Nova scheduler + testbed emulator (pymoc)
  - Synthetic and real datacenter trace-driven workload
  - Estimate performance degradation through host congestion
  - Compare with theoretical upper-bound "oracle" scheduler



## Pulsar high-level functioning

- **IBM Adaptive Utilization Accelerator for Virtualized Environments (PULSAR):**
  - Simple VM idleless detector (CPU utilization threshold)
  - Claim resources from idle VMs by throttling them (reducing their resource reservation, **cgroups in KVM**)
  - Use **adjusted capacity** (considering throttling) to provision and place more VMs in the system

© 2014 IBM Corporation

© 2014 IBM Corporation

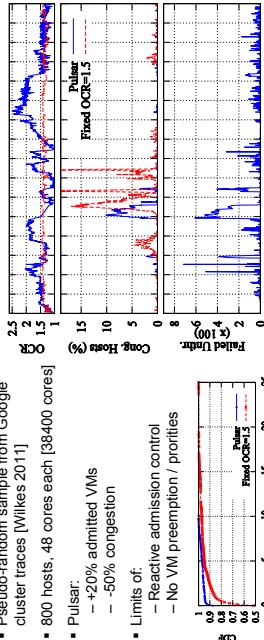
© 2014 IBM Corporation

## IBM Adaptive Utilization Accelerator for Virtualized Environments (PULSAR)



### Trace-based simulation

- Pseudo-random sample from Google cluster traces [Wikes 2011]
- 800 hosts, 48 cores each [38400 cores]
- Pulsar:
  - >20% admitted VMs
  - >-50% congestion
  - Limits of:
    - Reactive admission control
    - No VM preemption / priorities
- Fixed OCR=1.5



## IBM Adaptive Utilization Accelerator for Virtualized Environments (PULSAR)

### Conclusion

#### From our evaluation:

- PULSAR is adaptive to changes in resource utilization
- It increases infrastructure utilization
  - Limited host congestion
  - Limited number of VM migrations
  - Outperforming any fixed-OCR solution
- Future work:**
  - Use improved idleness detector / load predictors
  - Proactive admission control / VM priorities - preemption
  - Larger experiments!

#### ▪ Questions?

## IBM Adaptive Utilization Accelerator for Virtualized Environments (PULSAR)



### Experimental evaluation

- Trace-driven experiment (trace generated with boulder/sand model)
  - Daytrader (DT) Web app [3V/CPU, 30min period]
  - Sudoku solver (SD) [1V/CPU, 90min avg lifetime, 5% probability of switching b/w idle/active each minute]
  - Very "active" workload, very low maximum achievable OCR [1.5 max from Oracle]
- Tested
  - Openstack Controller node [Supermicro Xeon E5420 2.5 GHz cores, 8GB RAM]
  - 2 Openstack Compute nodes [IBM System X3550 M3, 24 Xeon X5680 3.3 GHz cores, 28GB RAM]
- Runs (averaged over 20 executions):
  - R1: 4 DTs + 36 SDs (group A), OCR=1
  - R2: PULSAR with group A + SDs from a Poisson process with 2 minutes inter-arrival time (group B)
  - R3: fixed OCR=1.27 (average obtained by Pulsar) groups A+B

Run	DT avg RT (std) [ms]	SD avg thr (STD) [Hz]	Host 1 avg util [%]	Host 2 avg util [%]
R1	28.8 (6.4)	61.2 (13)	57.8	62.3
R2	34.6 (9.7)	50.25 (14.1)	79.1	80.4
R3	41.6 (11.8)	46.95 (12.85)	85	84

© 2014 IBM Corporation

## IBM Adaptive Utilization Accelerator for Virtualized Environments (PULSAR)

### Backup slides



## Related work

### References

- [Wilkes 2011]: John Wilkes, "More Google cluster data", Google research blog, Nov 2011
- [Breitgand 2012] D. Breitgand and A. Epstein, "Improving Consolidation of Virtual Machines with Risk-Aware Bandwidth Oversubscription in Compute Clouds," in INFOCOM, 2012.
- [Chen 2011] M. Chen, H. Zhang, Y.-Y. Su, X. Weng, G. Jiang, and K. Yoshihira, "Effective VM Sizing in Virtualized Data Centers," in IEEE/IFIP IM'11, Dublin, Ireland, May 2011.
- [Meng 2010] X. Meng, C. Isci, J. Kephart, L. Zhang, E. Boulier, and D. Pandurakis, "Efficient Resource Provisioning in Compute Clouds via VM Multiplexing," in The IEEE/ACM International Conference on Autonomic Computing and Communications, Washington, DC, USA, Jul 2010.
- [Gnach 2007] D. Gnach, J. Rolia, L. Cherkasova, and A. Kemper, "Workload Analysis and Demand Prediction of Enterprise Data Center Applications," 2007 IEEE 10th International Symposium on Workload Characterization, pp. 171–180, Sep. 2007.
- [Gnach 2012] D. Gnach, J. Rolia, and L. Cherkasova, "Selling T-shirts and Time Shares in the Cloud," Cloud and Grid Computing, 2012.
- [Yanagisawa 2013] H. Yanagisawa, T. Osgami, and R. Raymond, "Dependable Virtual Machine Allocation," in Infocom, 2013, pp. 653–661.
- [Blagoiurov 2013] S. Blagoiurov, D. Gnach, M. Anitt, Y. Chen, C. Hyser, and A. Fedorova, "Maximizing Server Utilization while Meeting Critical SLAs via Weight-Based Collocation Management," in IFIP/IEEE IM'13, 2013.
- [Wulini 2012] F. Wulini, R. Städler, and H. Lindgreen, "Dynamic resource allocation with management objectives: Implementation for an OpenStack cloud," in CNSM'12, 2012, pp. 309–316.

## Related work

### References

- Many papers on demand prediction for stable VM population [Breitgand 2012] [Chen 2011]
  - We consider dynamic VM population, discrepancy between nominal and actual resource usage, adaptive over-commit
- [Gmach 2013] assume static VM population and no overcommit model,
  - We use past VM demand patterns to predict future demand
  - We left out prediction of future demand on purpose assuming dynamic VM population, albeit we could leverage this information
- [Carrera 2012] aim at fair placement decision by using a model of expected performance given a resource allocation for each workload
- [Blagoiurov 2013] requires application performance monitoring instrumentation and knowledge of resource consumption profiles to classify applications as batch or interactive
- [Wulini 2012] use average resource utilization over a sliding window to implement different placement policies (e.g., consolidation). The same solution can be applied for adaptive overcommit. However, churn and high utilization variation cause number of required migrations to grow quickly

**DYNAMIC POWER MANAGEMENT IN DATA CENTERS:  
THEORY & PRACTICE****Mor Harchol-Balter, Computer Science Department,  
Carnegie Mellon University**

Energy costs for data centers continue to rise, already exceeding ten billion dollars yearly. Sadly much of this power is wasted. Servers are only busy 10-30% of the time, but they are often left on, while idle, utilizing 60% of more of peak power while in the idle state. The obvious solution is dynamic power management: turning servers off, or re-purposing them, when idle. The drawback is a prohibitive "setup cost" to get servers back on. The purpose of this talk is to understand the effect of the "setup cost" and whether dynamic power management makes sense.

We first turn to theory and study the effect of setup cost in an M/M/k queue. We present the first analysis of the M/M/k/setup queueing system. We do this by introducing a new technique for analyzing infinite, repeating, Markov chains, which we call Recursive Renewal Reward (RRR).

We then turn to implementation, where we implement and evaluate dynamic power management in a multi-tier data center with key-value store workload, reminiscent of Facebook or Amazon. We propose a new dynamic algorithm, AutoScale, which is ideally suited to the case of unpredictable, time-varying load, and we show that AutoScale dramatically reduces power in data centers.

Joint work with: Anshul Gandhi, Alan Scheller-Wolf, and Mike Kozuch

## Power is Expensive

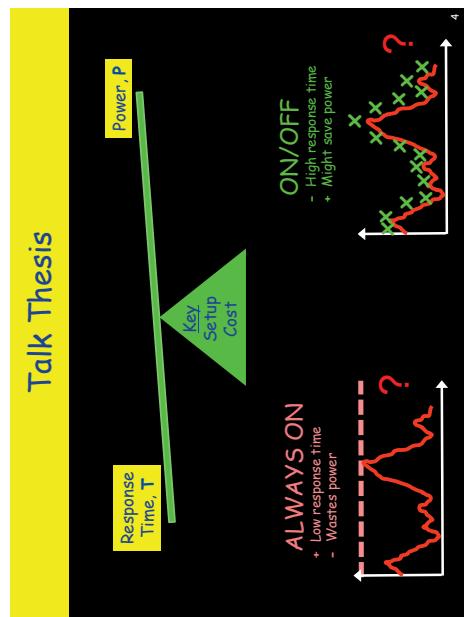
Annual U.S. data center energy consumption

- || 100 Billion kWh or 7.4 Billion dollars
- || Electricity consumed by 9 million homes
- || As much CO<sub>2</sub> as all of Argentina

[energystar.gov], [McKinsey & Co.], [Gartner]



2



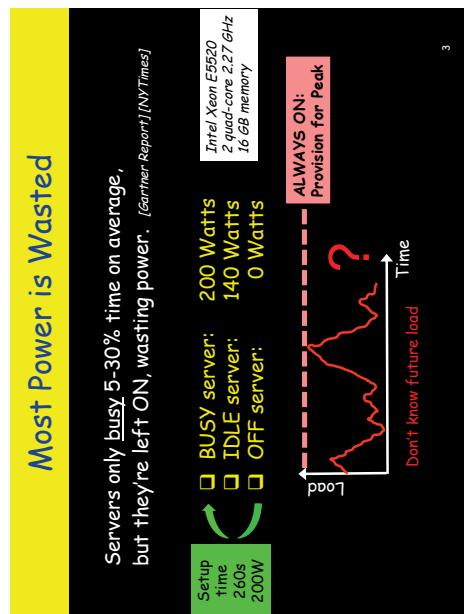
## Power Management in Data Centers: Theory & Practice

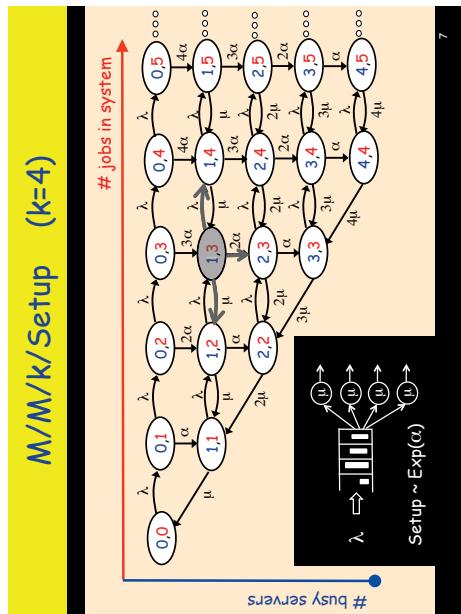
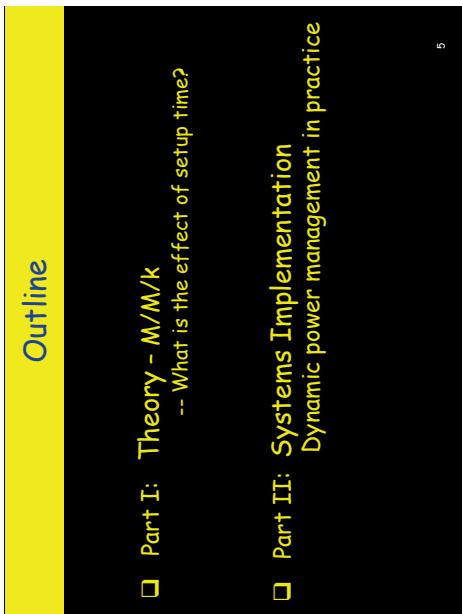
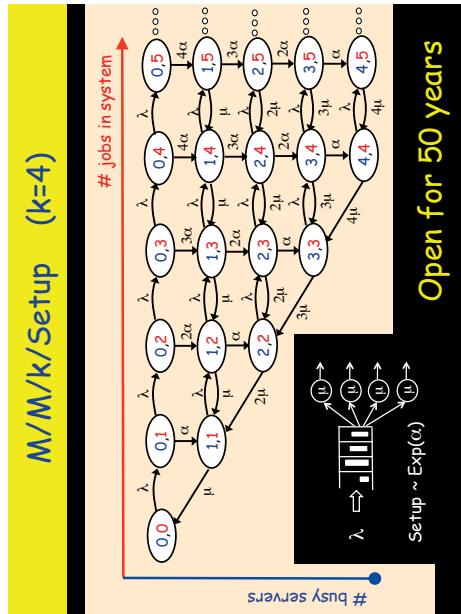
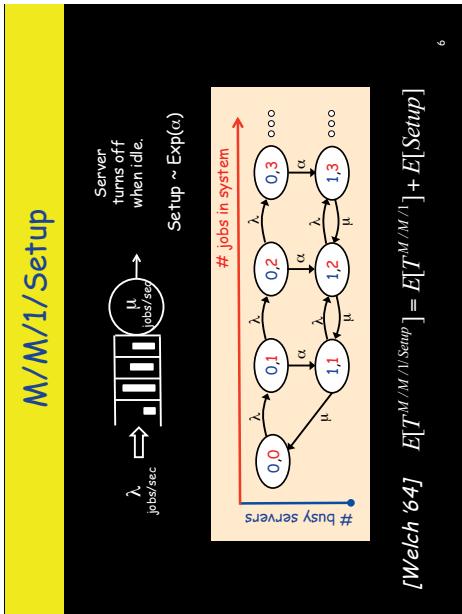
Mor Harchol-Balter  
Computer Science Dept  
Carnegie Mellon University

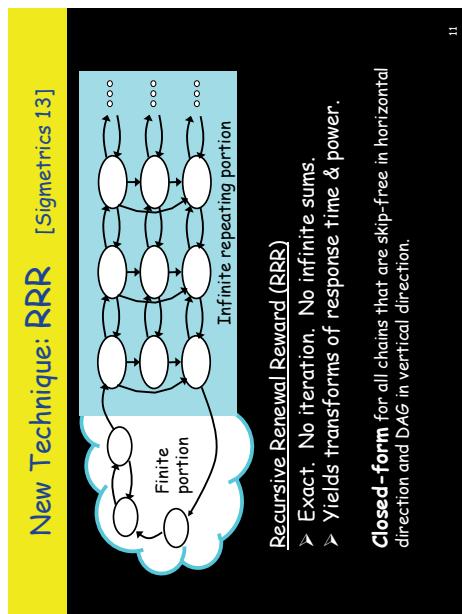
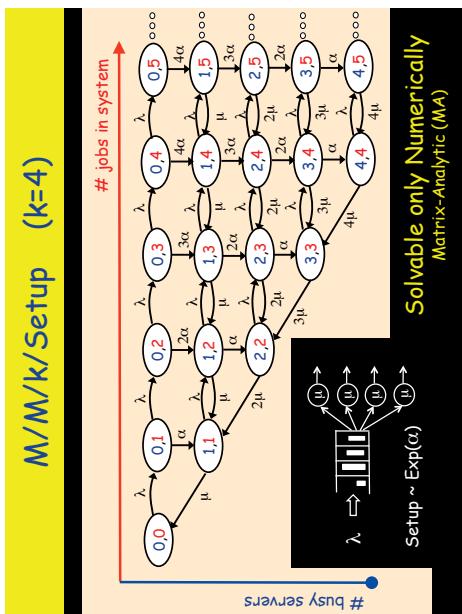
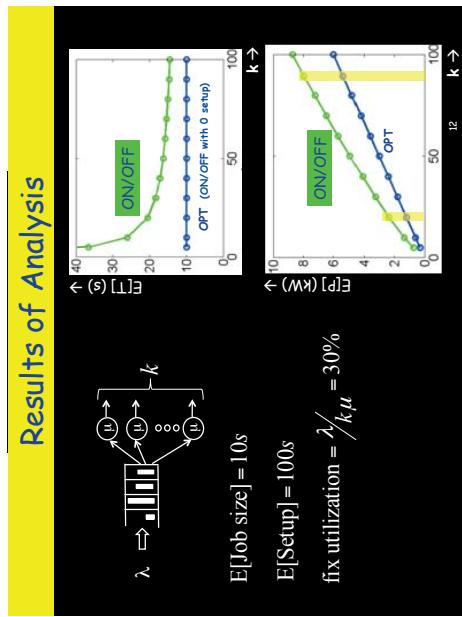
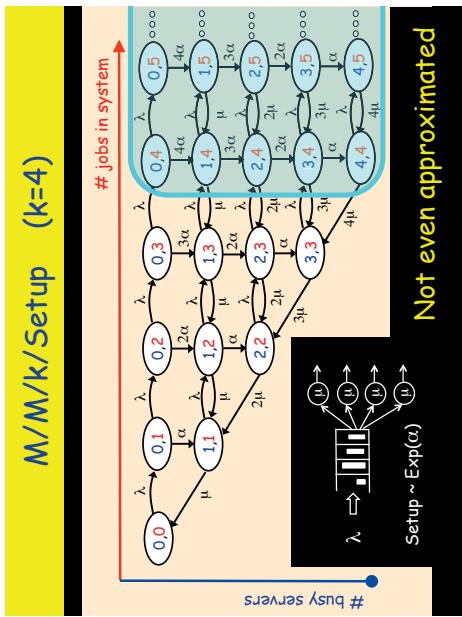
Anshul Gandhi, Sherwin Doroudi,  
Alan Scheller-Wolf, Mike Kozuch



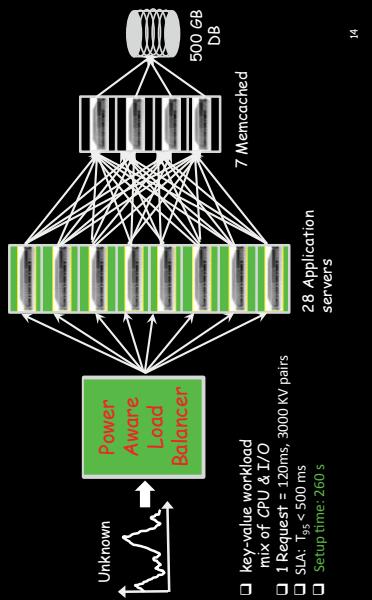
4







## Our Data Center

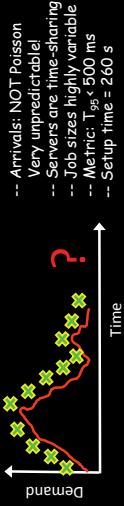


## Outline

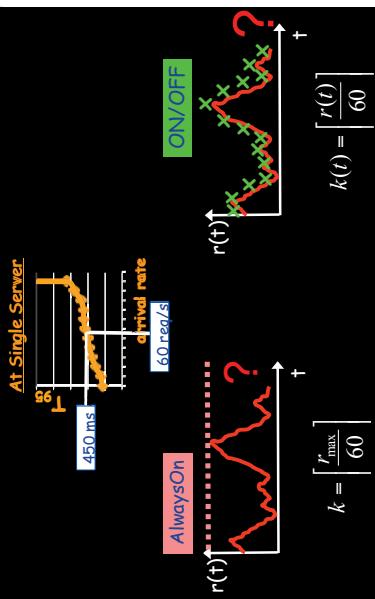
- Part I: Theory - M/M/k  
What is the effect of setup time?

-- Setup hurts a lot when k: small  
-- But setup much less painful when k: large  
-- ON/OFF allows us to achieve near-optimal power

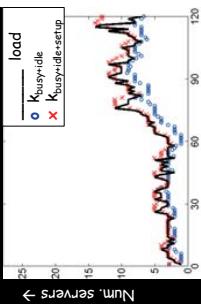
- Part II: Systems Implementation  
Dynamic power management in practice

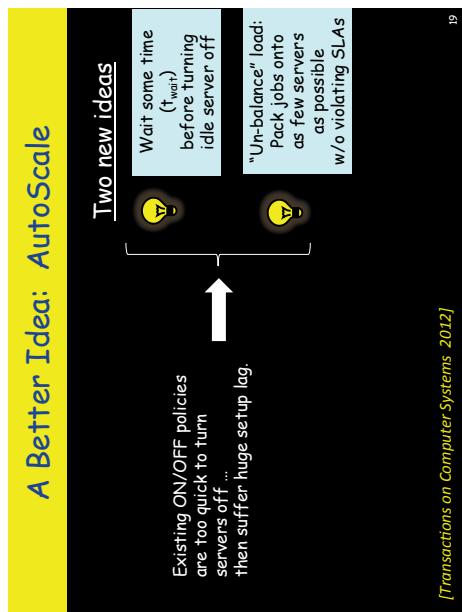
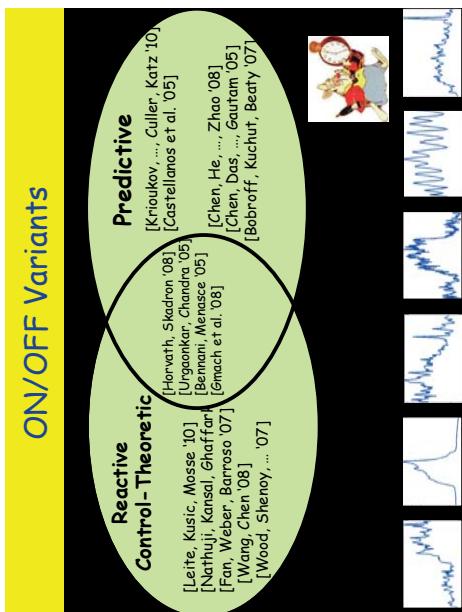
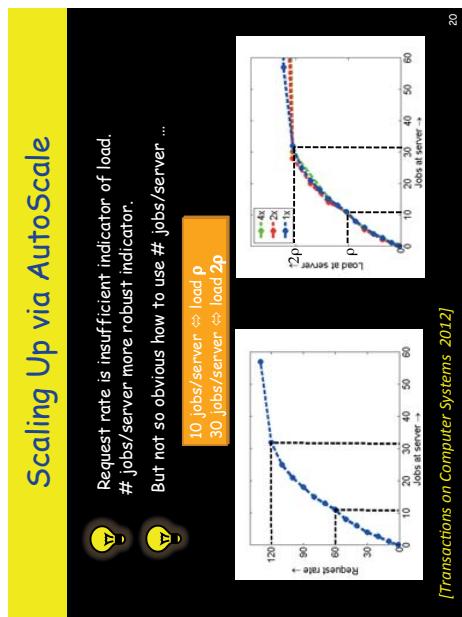
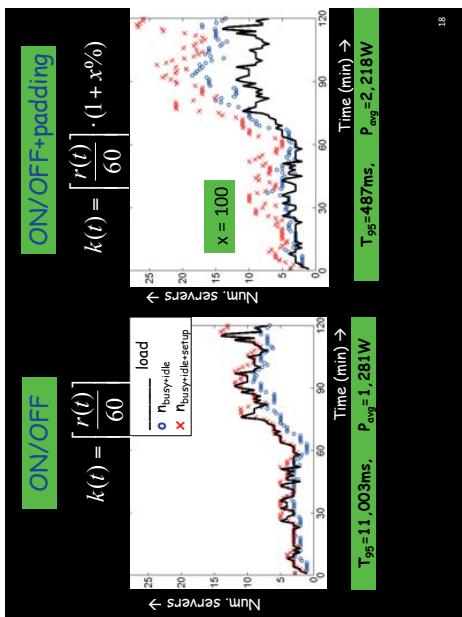


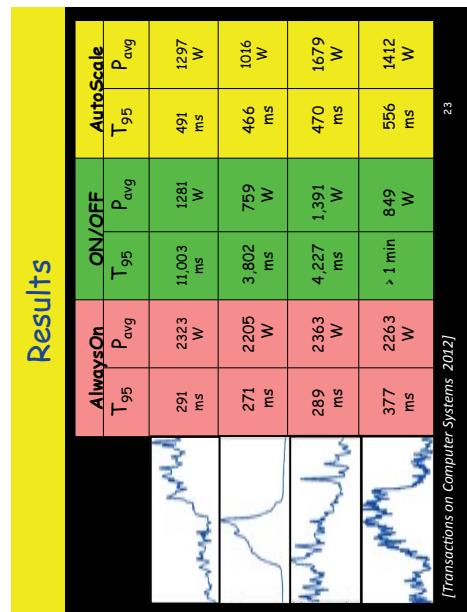
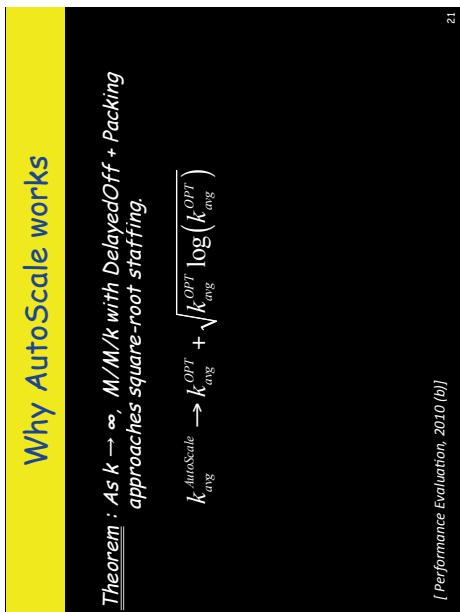
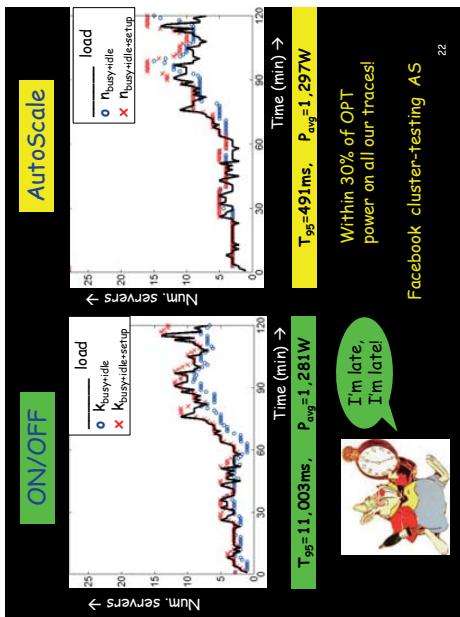
## Provisioning



## ON/OFF

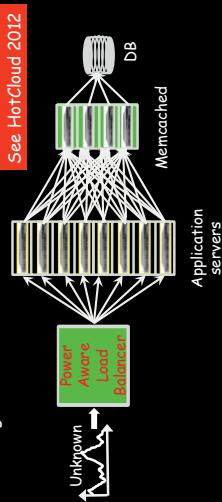






## Comments related to LCCC

□ Scaling stateful servers?



□ Tradeoffs between architectures:  
“Should we separate stateful from stateless?”

**See HotCloud 2012 - best of both**

25

## References

- Anshul Gandhi, Sherwin Doroudi, Mon Harchol-Balter, Alan Scheller-Wolf, "Exact Analysis of the M/M/k setup Class of Markov Chains via Recursive Renewal Reward," *ACM SIGMETRICS 2013 Conference*, June 2013.
- Anshul Gandhi, Mon Harchol-Balter, R. Rajgundaran, Mike Kozuch, "AutoScale: Dynamic, Robust Capacity Management for Multi-tier Data Centers," *ACM Transactions on Computer Systems*, vol 30, No. 4, Article 14, 2012, pp.1-26.
- Anshul Gandhi, Timothy Zhu, Mon Harchol-Balter, Mike Kozuch, "SOFTScale: Stealing Opportunistically For Transient Scaling," *Middleware 2012*.
- Timothy Zhu, Anshul Gandhi, Mon Harchol-Balter, Mike Kozuch, "Saving Cash by Using Less Cache," *HotCloud 2012*.
- Anshul Gandhi, Mon Harchol-Balter, and Ivo Adan, "Server farms with setup costs," *Performance Evaluation*, vol. 67, no. 11, 2010, pp. 1123-1138,
- Anshul Gandhi, Varun Gupta, Mon Harchol-Balter, and Michael Kozuch, "Optimality Analysis of Energy-Performance Trade-off for Server Farm Management," *Performance Evaluation* vol. 67, no. 11, 2010, pp. 1155-1171.

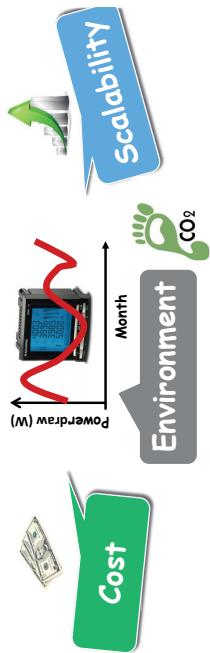
26

**REDUCING POWER DELIVERY COSTS FOR CLOUD DATA CENTERS****Bhuvan Urgaonkar, Penn State University**

Power-related costs are significant and growing components of overall data center expenditure (both capital and operational). In this talk, I will provide a brief overview of the sources and complexities of these costs, and of our work on using batteries and several IT knobs for reducing these costs. I will then provide a more detailed description of our use of stochastic control techniques for optimizing the data center's monthly electric utility bill. Finally, I will discuss complementary related work, open problems, and future directions.



## Data Centers Consume Lots of Power!



If treated as a country, fifth in the world for electricity use  
double in next 5 years, imposing a peak load of over 20 GW on the grid

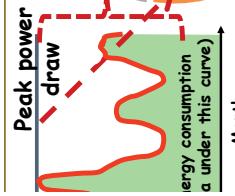
## Optimizing Peak Power-Related Costs in Cloud Data Centers

Bhuvan Urgaonkar  
Department of Computer Science and Engineering  
Penn State University

Collaborators: C. Wang, G. Kesidis



## Monthly Cost of 10MW Data Center



## Provisioned Peak Power Impact on Cap-ex

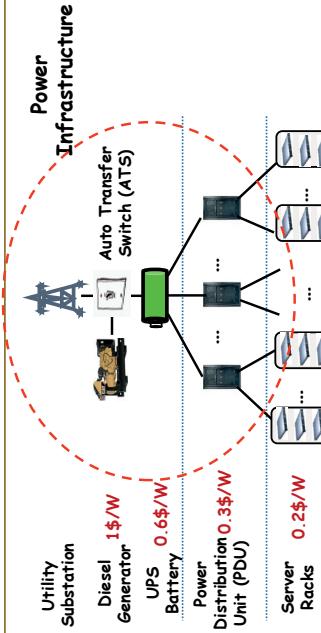
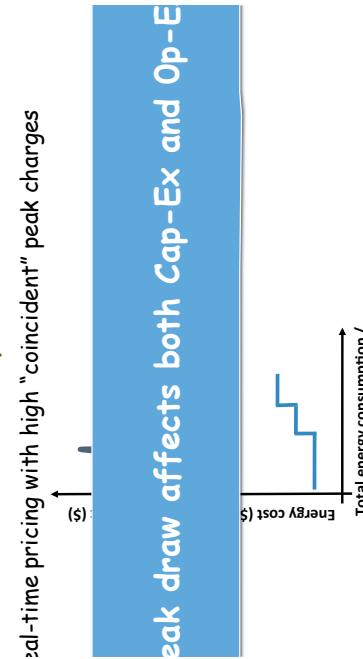
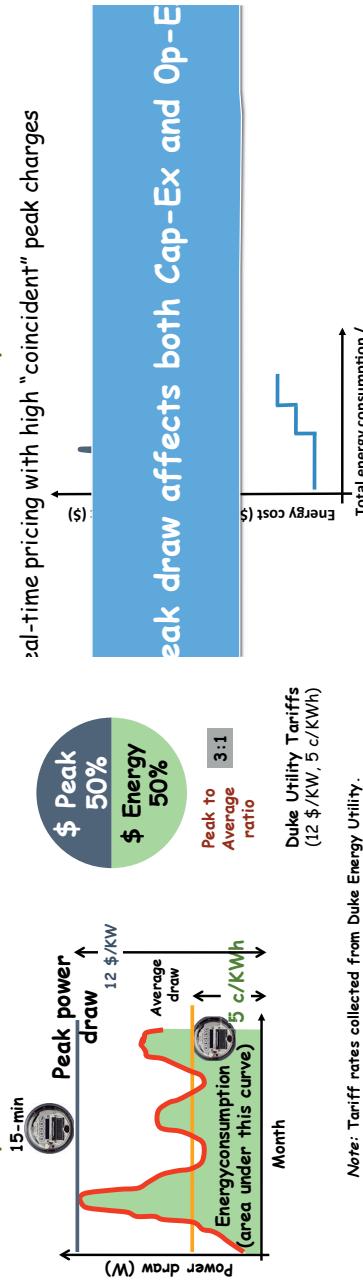


Chart:  
e. Book by Barroso et al., 1.5 pUE,  
-ap-ex, 20,000 servers, 1.5 pUE,  
ver & 12 yr infrastructure  
mortalization (Tier-2)  
*All cost are amortized  
at a monthly granularity*

## Consumed Peak Draw Contribution to Op-Ex (Explicit Peak-based Tariff)



## Optimizing Cap-Ex and Op-Ex

Cap-Ex optimization: How much capacity to provision for the Cap-Ex optimization: How much capacity to provision for the next several years?

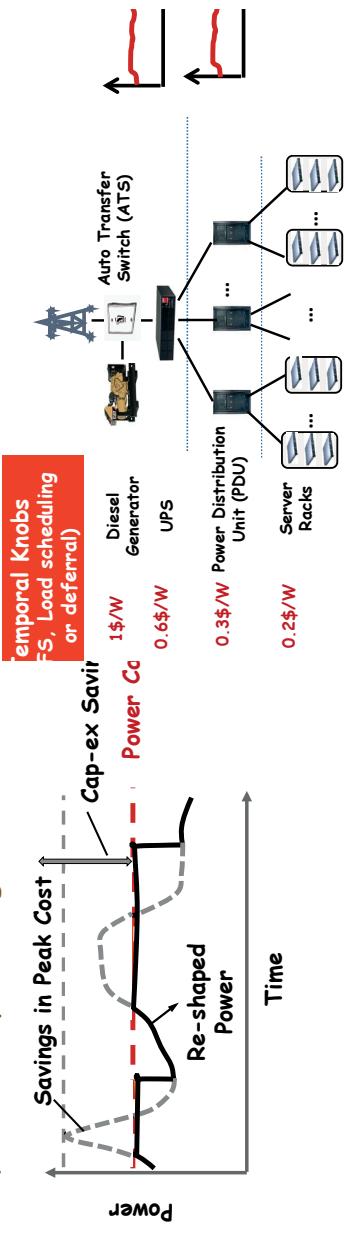
- An offline problem

Op-ex: How much peak to admit for **this billing cycle**?

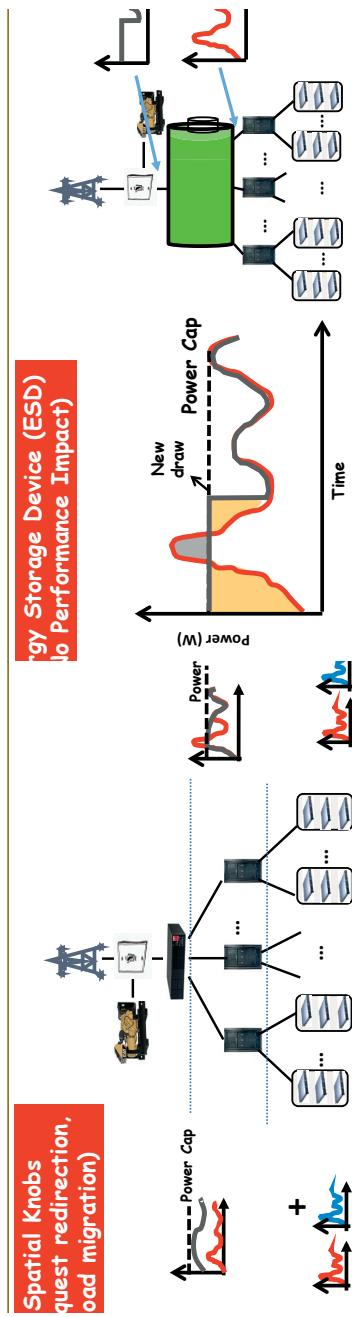
- An online control problem
  - \* Control windows may be in the minutes (or even seconds)

- Complementary problem: how to operate cost-effectively within a specified power capacity (as determined by cap-ex optimization)

## Demand Response: An Important Set of Techniques for Optimizing Power Costs



## Demand Response Knobs in a Data Center



## Overview of our Work

This talk

- Op-ex optimization using IT control knobs for a peak-based utility pricing scheme

Other work (happy to discuss offline)

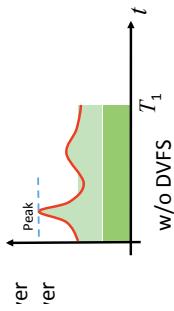
- Cap-ex improvements via provisioning of batteries and local generation sources
- Op-ex optimization:

- Real-time utility pricing schemes
- Control of batteries and local generation sources

## A Simple Model for IT-based DR

Despite their diversity, IT knobs can be viewed 'opping and/or delaying some power demand at the c performance degradation / revenue loss.

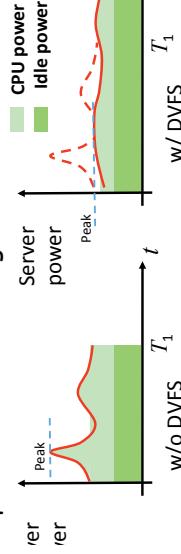
imple: DVFS/Scheduling



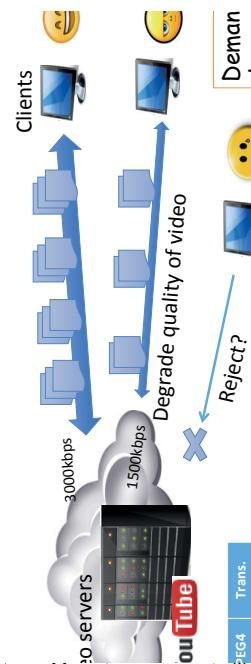
## A Simple Model for IT-based DR

Despite their diversity, IT knobs can be viewed 'opping and/or delaying some power demand at the c performance degradation / revenue loss.

imple: DVFS/Scheduling



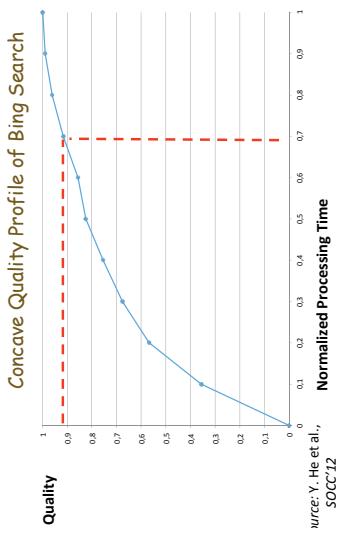
## Example 1: MPEG Video Server



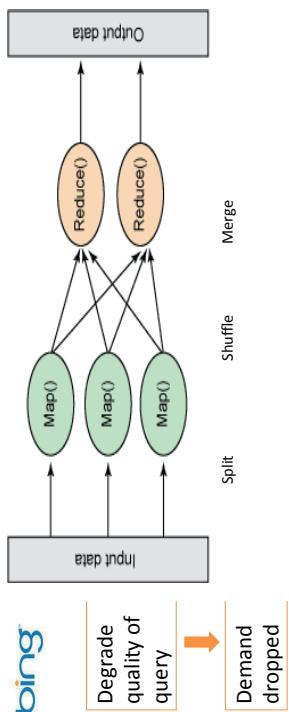
E.g., Source on revenue impact: "Video site quality impacts viewer behavior," Krishna Saraman, IEEE/ACM TON, 2013

Source: Y. Sharab et al., MMsys'13 and Wikipedia

## Example 2: Search Engine



## Example 3: Delay-tolerant, Batch



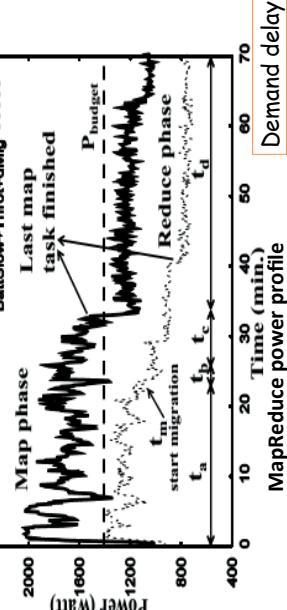
**bing**

Degrade quality of query

Demand dropped

## Example 3: Delay-tolerant, Batch

### Op-ex Optimization Problem



Source: S. Govindan et al., ASPLOS'12

How to use IT-based dropping or delaying of power demand to optimize op-ex vs. performance/revenue loss trade-off?



## Much Related Work for Real-time Pricing

### Real-time pricing

	Adversarial power demands	Stochastically known power demands		Adversarial power demands	Stochastically known power demands
sing IT-based DR	Z. Liu et al., Sigmetrics'13, robust optimization, avoid coincident peak		sing IT-based DR	Z. Liu et al., Sigmetrics'13, robust optimization, avoid coincident peak	<b>Current work (SDP formulation)</b>
sing batteries	R. Usgaonkar et al., Sig'11, Lyapunov optimization, distance from optimal inversely prop. to battery size P. Van de Ven et al., Energy'11, residential energy storage, MDP			R. Usgaonkar et al., Sig'11, Lyapunov optimization, distance from optimal inversely prop. to battery size P. Van de Ven et al., Energy'11, residential energy storage, MDP	

22

21

## But Less for Peak-based Pricing

### k-based pricing

	Adversarial power demands	Stochastically known power demands		Adversarial power demands	Stochastically known power demands
sing IT-based DR	Current work: CR=2 for time-varying energy prices; CR=2-1/T for fixed energy prices	Current work (SDP formulation, gSBB heuristic)	sing IT-based DR	Current work: CR=2 for time-varying energy prices; CR=2-1/T for fixed energy prices	Current work (SDP formulation, gSBB heuristic)
sing batteries for DR	A. Bar-Noy et al., WEA'08, threshold-based, CR of $H_n (=7.84$ if 30-min time-slot)	Current work (SDP formulation, gSBB heuristic)	sing batteries for DR	A. Bar-Noy et al., WEA'08, threshold-based, CR of $H_n (=7.84$ if 30-min time-slot)	Current work (SDP formulation, gSBB heuristic)

24

23

## But Less for Peak-based Pricing

### k-based pricing

	Adversarial power demands		Adversarial power demands	Stochastically known power demands
sing IT-based DR	Current work: CR=2 for time-varying energy prices; CR=2-1/T for fixed energy prices		sing IT-based DR	Current work (SDP formulation, gSBB heuristic)
sing batteries for DR	A. Bar-Noy et al., WEA'08, threshold-based, CR of $H_n (=7.84$ if 30-min time-slot)		sing batteries for DR	Current work (SDP formulation, gSBB heuristic)

24

## Online "Dropping" of Power Demand

### Demand Response to Optimize Peak-based Utility Bill

Let's begin by assuming that the "knob" available to the data center is that of dropping part of the power demand

- Dropped demand never returns

Recall examples of a video streaming server and a search engine

**How to determine the peak demand to admit in an *online* fashion?**



## Offline Formulation for Dropping Demand

**Demand dropping**

- $l_{drop}(x)$ : Dropping demand v.s. Performance/Revenue loss
- Discretized optimization horizon  $T$ : A billing cycle (typically a month)
- Known demand time series  $\{p_t\}_{t=1}^T$

$$\min_{\{A_t\}, \{D_t\}} \beta y_{max} + \sum_{t=1}^T \{\alpha_t a_t + l_{drop}(d_t)\}$$

Peak cost

Dropping cost

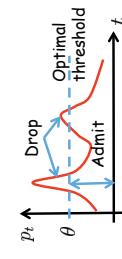
Energy cost

$$\begin{aligned} & s.t. \quad p_t - a_t - d_t = 0, \forall t \\ & \quad y_{max} \geq a_t, \forall t \\ & \quad \text{Peak of admitted demand} \end{aligned}$$

## Online Control: $ON_{drop}$

No information about future demand

Peak charge + time-varying energy price



## Online Control: $ON_{\text{drop}}$

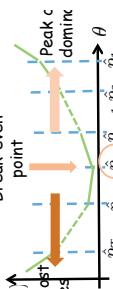
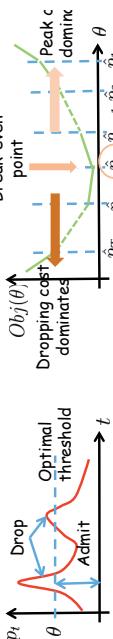
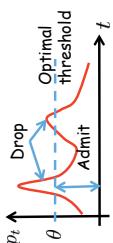
No information about future demand  
Peak charge + time-varying energy price

$$l_{\text{drop}}(x) = k_{\text{drop},x}$$

No information about future demand  
Peak charge + time-varying energy price

$$l_{\text{drop}}(x) = k_{\text{drop},x}$$

**Lemma.** The optimal solution has a demand dropping threshold  $\theta$  of the following form: if we denote as  $\hat{p}_t$  the  $t$ -th largest demand value in  $\{\hat{p}_t\}_{t=1}^T$  and as  $\hat{\alpha}_t$  the corresponding energy price, then  $\theta = \hat{p}_n$  and  $\beta - \sum_{t=1}^{n-1} (k_{\text{drop}} - \hat{\alpha}_t) \geq 0$ .



## Online Control: $ON_{\text{drop}}$

### Stochastic Control for Dropping Demand

In many cases, workloads can be predicted  
- Often via Markovian models

Can develop a SDP that leverages such predictive models

Offline formulation:  

$$\min_{\theta} E \left\{ \beta y_{\max} + \sum_{t=1}^T (\alpha_t a_t + l_{\text{drop}}(d_t)) \right\}$$

Stochastic dynamic programming?  

$$\begin{array}{c} \text{Sum} \\ \downarrow \\ \text{Sum + Max} \end{array}$$

Unconventional state space  
due to sum + max

**Theorem.**  $ON_{\text{drop}}$  offers a competitive ratio of 2 under peak-based pricing.

$$y_{t+1} = \max\{y_t, \alpha_t\}$$

## Stochastic Control for Dropping Demand Making the model a bit more complex

SDP<sub>Drop</sub> optimality rules:

$$V_t(y_t, p_{[t-1]}, \alpha_{[t-1]}) = \min_{\{A_t\}, \{D_t\}} E \{ \alpha_t a_t + l_{drop}(d_t) + V_{t+1}(y_{t+1}, p_t, \alpha_t) \}$$

$$\mid P_{[t-1]} = p_{[t-1]}, \Lambda_{[t-1]} = \alpha_{[t-1]} \}$$

$$\begin{aligned} s.t. \quad & y_{t+1} = \max\{y_t, a_t\} \\ & p_t - a_t - d_t = 0 \end{aligned}$$

**Lemma.** Under stage-independent demand SDP<sub>Drop</sub> has the following threshold-based optimal control policy :

$$(a_t^*, d_t^*) = \begin{cases} (\phi_t p_t, p_t - \phi_t p_t), & \text{if } \phi_t \leq 1 \\ (p_t, 0), & \text{if } \phi_t > 1 \end{cases}$$

What if dropping alone does not capture DR behavior?



Recall example of MapReduce ...

## Offline Problem Formulation

Demand delaying  $l_{delay}(x, t)$ : Delay up to  $\tau$  time slots  
Demand dropping  $l_{drop}(x)$ : Energy cost

$$\min_{\{\alpha_t\}, \{D_t\}} \beta y_{max} + \sum_t \{ \alpha_t (\sum_{i \in h^+(t)} a_{i,t}) + \sum_{i \in h(t)} l_{drop}(d_{i,t}) + \sum_{i \in h(t)} l_{delay}(a_{i,t}, t -$$

$$\begin{aligned} s.t. \quad & p_t - a_{i,t} - d_{i,t} = r_{i,t+1}, \forall i \\ & r_{i,t} - a_{i,t} - d_{i,t} = r_{i,t+1}, i \in h(t), \forall t \\ & r_{t-\tau,t} - a_{t-\tau,t} - d_{t-\tau,t} = 0, \forall t \\ & r_{i,T+1} = 0, i \in h(t) \\ & y_{max} \geq \sum_{i \in h^+(t)} a_{i,t}, \forall t \end{aligned}$$

New demand either admitted or dropped

$$r_{i,t} - a_{i,t} - d_{i,t} = r_{i,t+1}, i \in h(t), \forall t$$

Pending demand either admitted or dropped

$$r_{t-\tau,t} - a_{t-\tau,t} - d_{t-\tau,t} = 0, \forall t$$

Delayed for  $\tau$  time slots: Admit immediate

$$(s_t, p_{[t-1]}, \alpha_{[t-1]}) = \min_{\{A_t\}, \{D_t\}} E \{ \alpha_t a_t^+ + l_{drop}(d_{i,t}) + \sum_{i \in h(t)} l_{delay}(a_{i,t}, t -$$

No more delay at the end of billing cycle

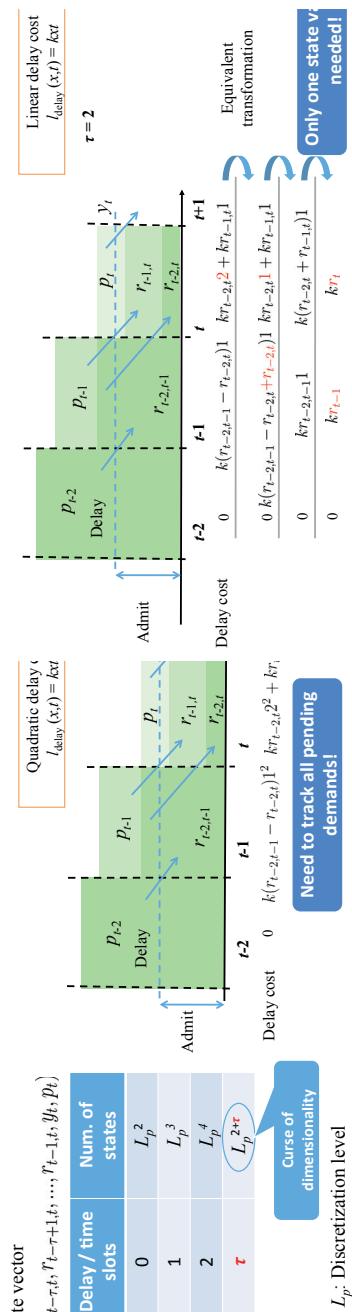
Peak of admitted demand

## Stochastic Control

SDP formulation  
- Track all pending demand if  $l_{delay}(x, t)$  is non-linear w.r.t.  $t$   
- Curse of dimensionality:  $O(RL_p^{2(\tau+2)L_a T})$

$$V_{t+1}(s_{t+1}, p_{[t]}, \alpha_{[t]}) \mid P_{[t-1]} = p_{[t-1]}, \Lambda_{[t-1]} = \alpha_{[t-1]})$$

## Stochastic Control: Curse of Dimensionality Stochastic Control: Linear delay cost



Equivalent transformation

## Scalable Approx. for SDP

- SDP<sub>Lin</sub>
- Linear approximation for  $I_{\text{delay}}$
  - $O(RL_p^{\frac{5}{2}}L_aT)$
- $$s_t = (r_t, y_t)$$
- $$V_t(s_t, p_{[t-1]}, \alpha_{[t-1]}) = \min_{\{A_t\}, \{D_t\}} E\{\alpha_t a_t + l_{\text{drop}}(d_t) + I_{\text{delay}}(r_t) + V_{t+1}(s_{t+1}, p_{[t]})\}$$
- $$\quad | \quad P_{[t-1]} = p_{[t-1]}, \Lambda_{[t-1]} = \alpha_{[t-1]}\}$$
- $$\quad s.t. \quad y_{t+1} = \max\{y_t, a_t\}$$
- $$\quad r_{t+1} = (p_t - d_t) - a_t - r_t$$

What if SDP does not scale?



A scalable technique based on a "GSBB" model for power demand

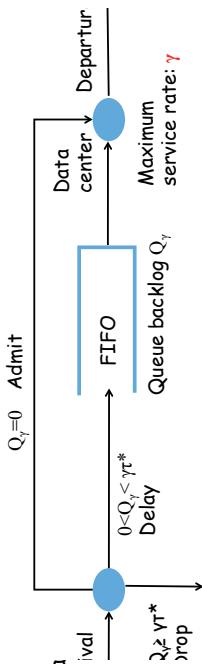
## qSBB-based Control

Raw demand is modeled as "generalized stochastically bounded" or "burstiness" curve

$$\{(\gamma, \phi(\gamma\tau^*)) \mid \gamma > \mu\}$$

A queue whose arrivals are the "raw" demands and is served at rate  $\gamma$  will have backlog  $Q_\gamma$  such that

$$Pr(Q_\gamma \geq \gamma\tau^*) \leq \phi(\gamma\tau^*)$$



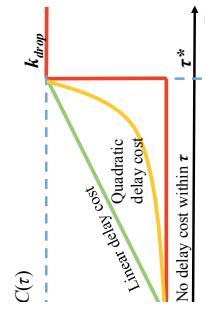
## qSBB-based Control

How to obtain  $\gamma$ ?

$$Pr(Q_\gamma \geq \gamma\tau^*) \leq \phi(\gamma\tau^*)$$

## qSBB-based Control

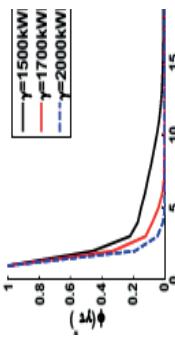
$$\text{Objective} \quad \min_{\gamma > \mu} T \mu \int_0^\infty C(\tau) dF_\gamma(\tau) + \beta \gamma - \phi(\gamma\tau)$$



Examples of  $C(\tau)$

## Selected Simulation Results

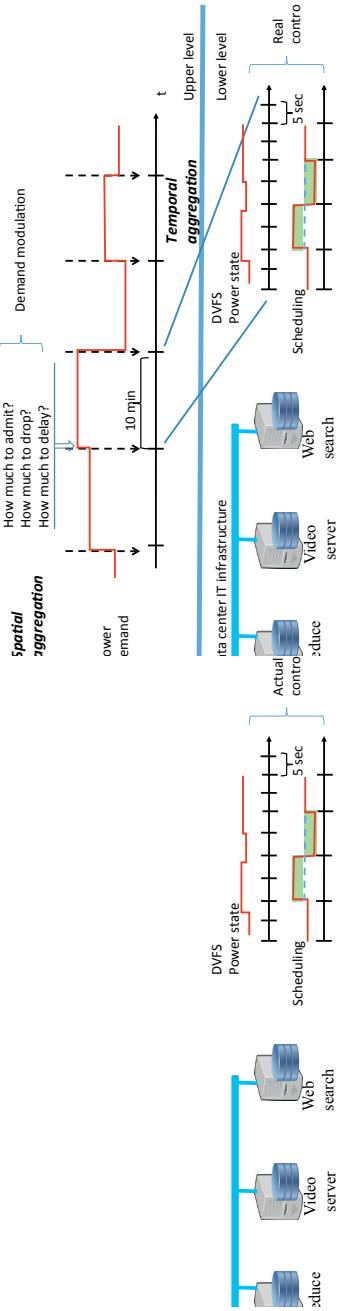
Cost benefits of demand response via abstract demand dropping and delaying  
From abstract control to real control: A case study



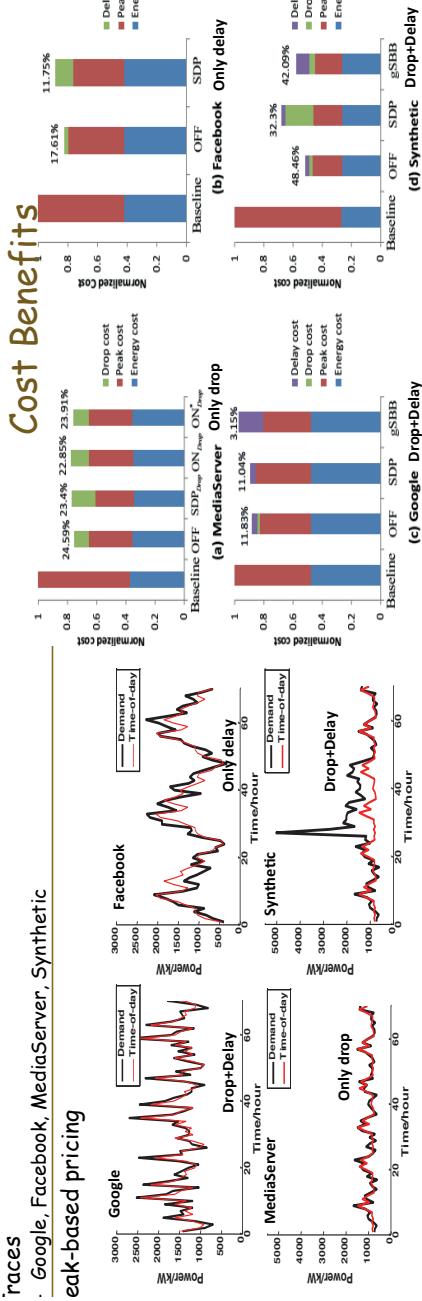
Example of  $\phi(\gamma\tau^*)$

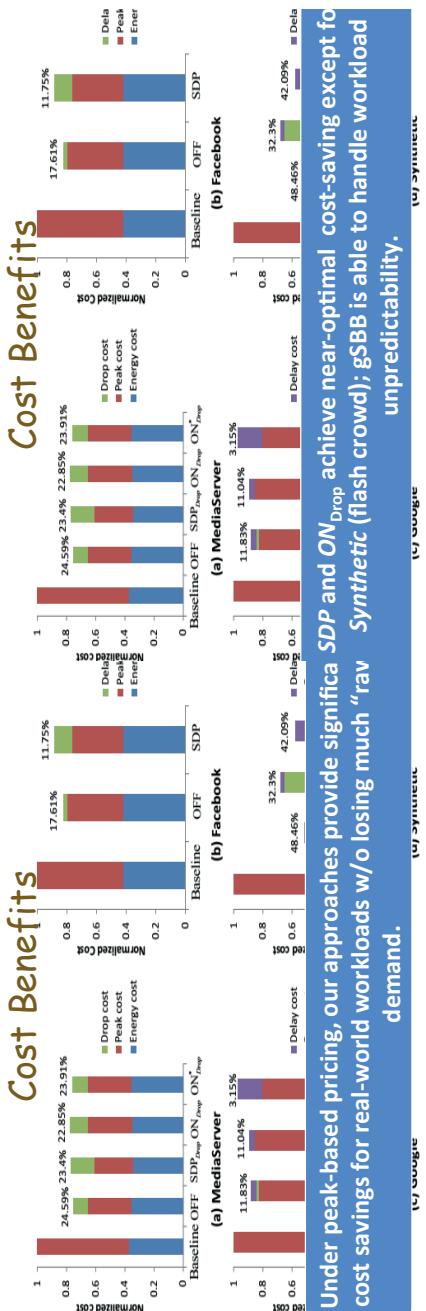
## A Hierarchical Demand Response Framework

### A Hierarchical Demand Response Framework



Figures 10: Google, Facebook, MediaServer, Synthetic  
peak-based pricing

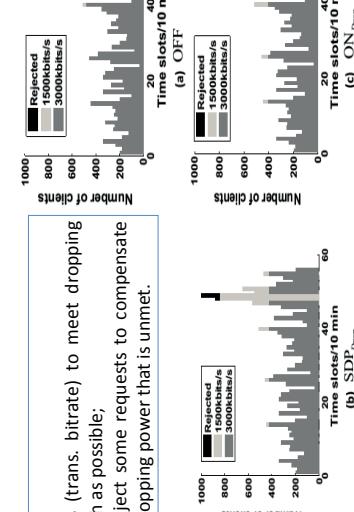




## A Case Study: Media Server

## Conclusions and Open Problems

- Peak power draw significantly impacts both cap-ex and op-ex
  - Algorithms and empirical case studies from our work on such optimization using IT knobs
  - Key idea: Abstract myriad IT knobs as dropping or delaying power demand at the cost of performance/revenue loss
  - Results for both adversarial inputs and stochastically known inputs
  - Plenty of scope for more work (both theoretical and empirical) on op-ex optimization for peak-based pricing schemes, e.g.:
    - Competitive analysis for real-time pricing using batteries for DR
    - Competitive analysis for peak-based pricing using IT knobs and/or batteries when using both "dropping" and "delaying" of power demand
- More details at: <http://www.cse.psu.edu/~bhuvan>



### Installation.

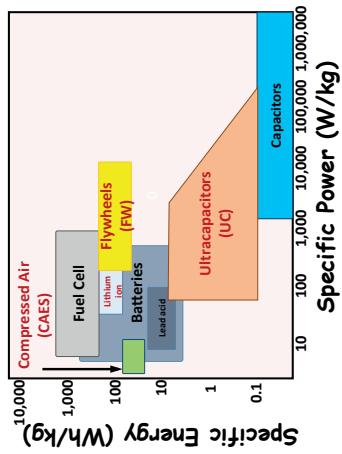
Degrade QoS (trans. bitrate) to meet dropping decision as much as possible;  
otherwise, reject some requests to compensate the target dropping power that is unmet.

## ESDs in Current Datacenters

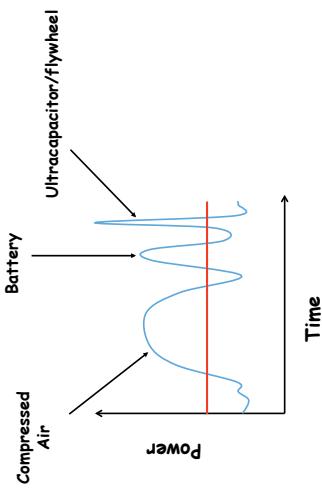
Why restrict ESDs to any one level of the datacenter power hierarchy (e.g., central or server)?

Why restrict to single ESD technology (e.g., Lead acid battery)?

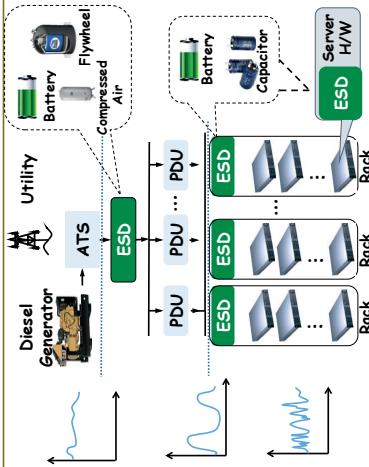
## Ragone Plot



## Hybrid ESD solution may be desirable



## Multi-level Multi-technology ESDs



## BROWNOUT: BUILDING MORE ROBUST CLOUD APPLICATIONS

Cristian Klein, Umeå University

Resource allocation in clouds is mostly done assuming hard requirements, applications either receive the requested resources or fail. Given the dynamic nature of workloads and the risk of cascading failures, guaranteeing on-demand allocations requires large spare capacity. Hence, one cannot have a system that is both reliable and efficient.

To solve this issue, we introduce brownout, a new paradigm to improve the robustness of replicated cloud applications. Brownout applications contain some optional code that can be dynamically deactivated as needed. Although this idea is simple and fairly non-intrusive to application code, properly supporting it required changes in several components. First, at the replica level, we synthesize a replica controller to decide when to execute the optional code and when to skip it. Second, we propose a resource manager to decide allocations among multiple brownout applications in a fair manner. Third, we propose two novel load-balancing algorithms, specifically designed for brownout replicas, to maximize the amount of optional content served. We theoretically prove properties of the overall system using control and game theory.

To show the practical applicability, we implemented brownout versions of RUBiS and RUBBoS with less than 170 lines of code. The load-balancing algorithms were implemented on top of lighttpd with less than 180 lines of code. Experiments show that brownout may enable considerable improvements in withstanding flash-crowds or hardware failures. Brownout opens up more flexibility in cloud resource management, which is why we encourage further research by publishing all source code.





## Cloud Computing Definition

### How I Learned to Stop Worrying and Love Capacity Shortages

Cristian Klein  
Umeå University



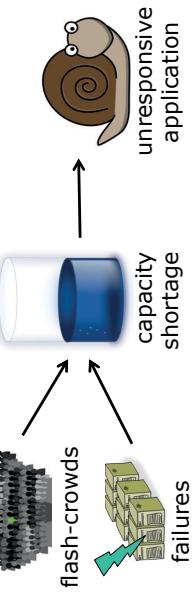
2014-05-08, LCCC Workshop on Cloud Control  
Lund, Sweden

**Rapid provisioning** (and release)  
from a **shared pool** of resources

- 3 deployment models
  - Public cloud ("our stuff")
  - **Private cloud** ("my stuff")
    - Hybrid cloud (combine the two above)
- 3 service models (what is leased)
  - Software-as-a-Service (SaaS)
  - Platform-as-a-Service (PaaS)
  - Infrastructure-as-a-Service (IaaS)

2

### Infrastructure-as-a-Service (IaaS)



- 82% of end-users give up on a lost payment transaction\*
- 25% of end-users leave if load time > 4s\*\*
- 1% reduced sale per 100ms load time\*\*
- 20% reduced income if 0.5s longer load time\*\*\*

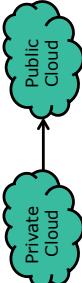
3

4

\* JupiterResearch   \*\* Amazon   \*\*\* Google

## State-of-Practice

- Large spare capacity
    - May be economically impractical
- 

- Cloud bursting
    - Lease capacity from a public cloud
    - Does not really solve the problem
- 

5

## Brownout: Idea

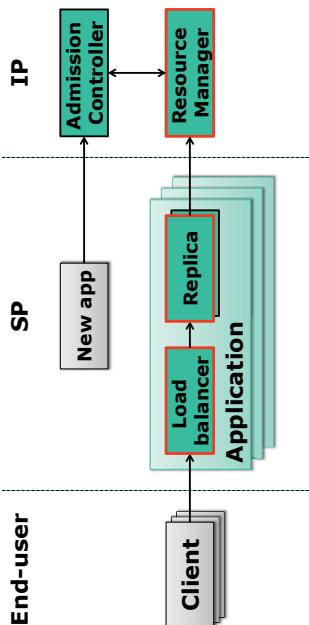
- Disable **optional** content
  - Minimally intrusive
  - E.g. recommendations
  - ~50% increase in sales \*
- Challenge
  - Maximize optional content
  - Avoid high response times



6

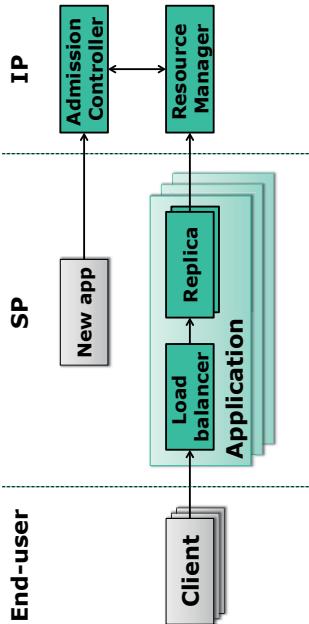
\* D. Feder et al., "Recommendation systems and their effects on consumers: the fragmentation debate," in Electronic Commerce, 2010. DOI: 10.1145/1807342.1807378.

## Cloud Architecture



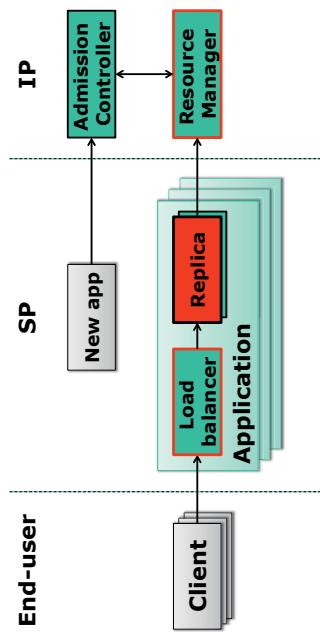
7

## Cloud Architecture



8

## Cloud Architecture

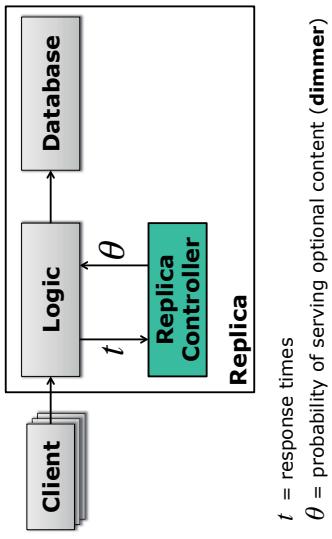


9

10

C. Klein, M. Maggio, F. Hernández-Rodríguez, K-E. Ågren,  
"Brownout: building more robust cloud applications", ICSE, 2014

## Brownout: Inside a Replica



9

10

## Replica Controller (1)

- Need to adapt to changes
  - Number of users
  - Available capacity
- Not all requests take the same time
  - E.g., cached in memory, disk
- Need to reject disturbances
  - E.g., NTP daemon firing up, cron jobs

## Replica Controller (2)

- Start from a simple model

$$t^{k+1} = \alpha^k \cdot \Theta^k + \delta t^k$$

- Adaptive PI controller

$$\Theta^{k+1} = \Theta^k + \frac{1 - p_1}{\alpha} \cdot e^{k+1}$$

- $\alpha$  estimated using RLS

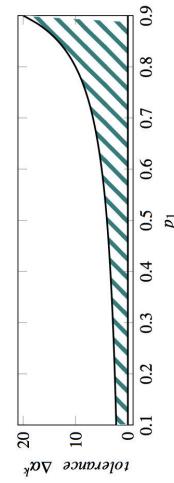
M. Maggio, C. Klein, K-E Ågren,  
"Control strategies for predictable brownouts in cloud computing", IFAC, 2014

11

12

## Robustness to Model Uncertainties

$$\alpha = \tilde{\alpha} \cdot \Delta\alpha$$

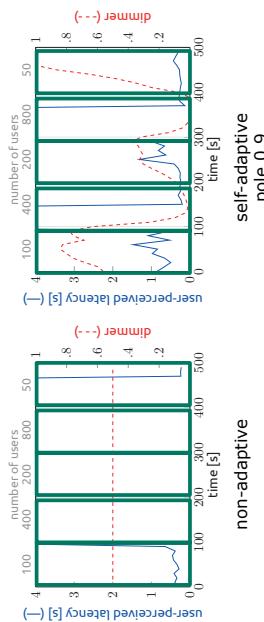


## Evaluation

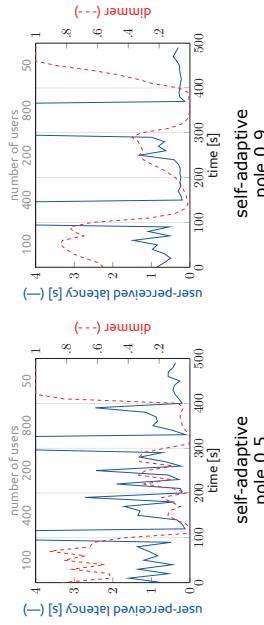
- RUBiS: eBay-like prototype
  - Added a recommender
- RUBBoS: Slashdot-like prototype
  - Added a recommender
  - Marked comments as optional
- Effort in lines of code:

Modification	RUBiS	RUBBoS
Recommender	37	22
Dimmer	3	6
Reporting response time to controller	5	5
Controller	120	120
<i>Total</i>	<i>165</i>	<i>153</i>

## Results: RUBiS, flash-crowd Non-adaptive vs. self-adaptive



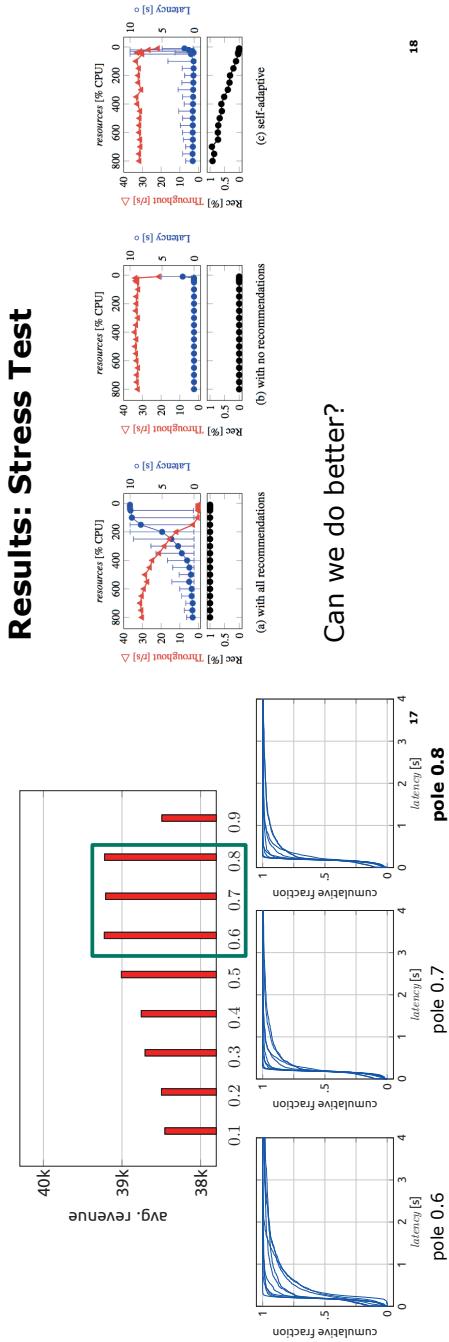
## Results: RUBiS, flash-crowd Self-adaptive: different pole values



15

16

**Results:** revenue = 1 × #requests + 0.5 × #optional



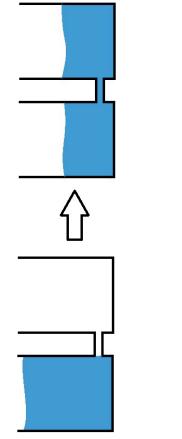
## Results: Stress Test

Can we do better?

## Cloud Architecture

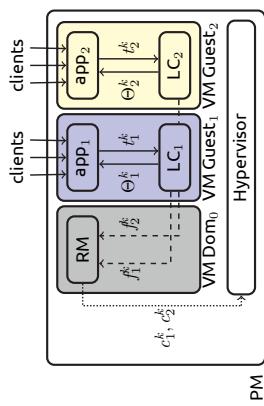
## Goal

- Give capacity to application that "struggles"
- **Fairly** balances capacity among applications



## Zoom: Inside a Physical Machine

### Details

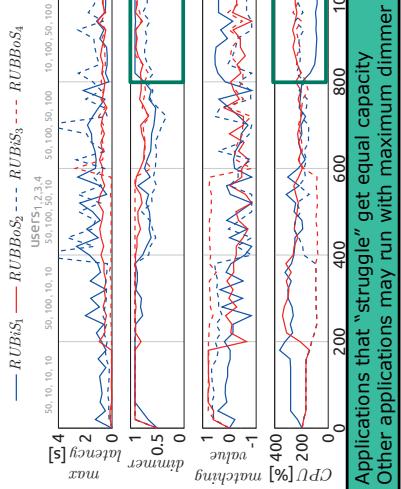


- Application sends **matching values**
  - Resource manager computes **capacities**
  - Proven to **converge** and be **fair** using game theory
- $$f_i^k = 1 - t_i^k / \bar{t}_i$$
- $$c_i^{k+1} = c_i^k - \varepsilon_{rm} \left( f_i^k - c_i^k \cdot \sum_p f_p^k \right)$$

C. Klein, M. Maggio, F. Hernández-Rodríguez, K-E Álvarez,  
"Resource management for service level aware cloud applications", REACTION, 2013

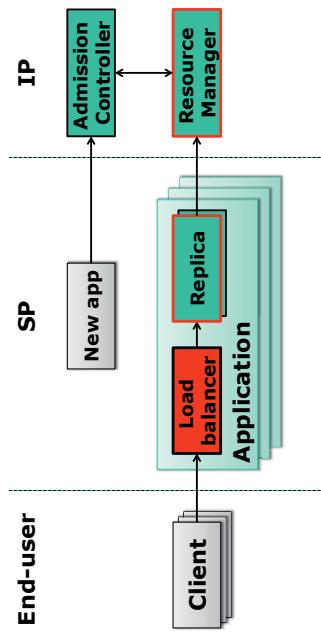
21

## Results: 4 Applications



Applications that "struggle" get equal capacity  
Other applications may run with maximum dimmer

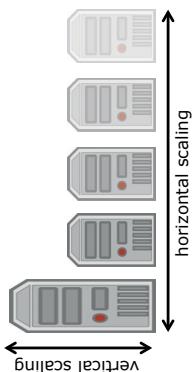
## Cloud Architecture



24

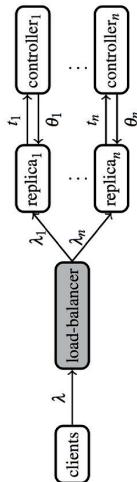
## Why Replication?

- Scale beyond a physical machine



## Why Replication and Brownout?

- Hide auto-scaling mishaps
- Hide failures leading to **capacity shortage**

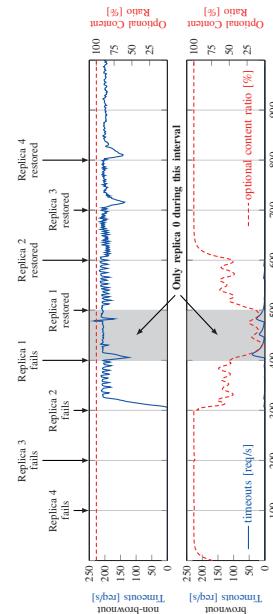


- Resilience, **hide** infrastructure failures

25

## Results: Replication and Brownout

**Brownout-unaware load-balancing: shortest queue first**



27

## Maximizing Optional Content

	Weight-based (periodic)	Queue-based (event)
Brownout-unaware		SQF
Variation-based	WRR	P1BH + EPBH +
Equality-based	PIBH	COBLB
Optimization-based	EPBH	

### Tested using simulations

- SQF best brownout-unaware method
- Brownout-aware methods better
  - Somewhat slow to react

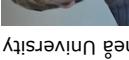
### Tested using experiments (lighttpd)

J. Durango et al.: "Control-theoretical load-balancing for cloud applications with brownout",  
 (submitted to CDC 2014)  
 C. Kien et al.: "Improving Cloud Service Resilience using Brownout-Aware Load-Balancing",  
 (submitted to SDOs 2014)

28

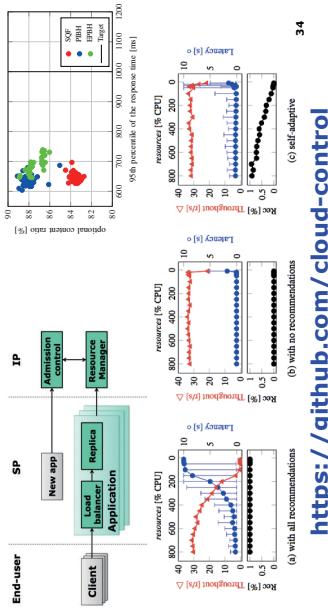
Thank you for your attention!

## Acknowledgments

	Luis Tomás, PhD, Post-doc
	Johan Tordsson, PhD, Assistant Professor
	Francisco Hernández, PhD, Assistant Professor
	Erik Elmroth, PhD, Professor
	Umeå University Lund University

Cristian Klein

## How I Learned to Stop Worrying and Love Capacity Shortages



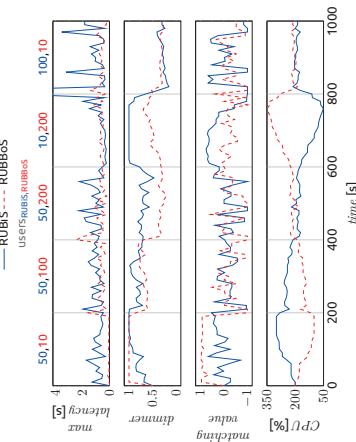
<https://github.com/cloud-control>

## Appendix

1. C. Klein, M. Maggio, F. Hernández-Rodríguez, K-E Årzén, "Brownout: building more robust cloud applications", ICSE, 2014
2. M. Maggio, C. Klein, K-E Årzén, "Control strategies for predictable brownouts in cloud computing", IFAC, 2014
3. C. Klein, M. Maggio, F. Hernández-Rodríguez, K-E Årzén, "Resource management for service level aware cloud applications", REACTION, 2013
4. J. Durango, M. Delkrantz, M. Maggio, C. Klein, A. V. Papadopoulos, F. Hernández-Rodríguez, E. Elmroth, "Control-theoretical load-balancing for cloud applications with brownout", (submitted to CDC 2014)
5. C. Klein, A. V. Papadopoulos, M. Delkrantz, J. Durango, M. Maggio, K-E Årzén, F. Hernández-Rodríguez, E. Elmroth, "Improving Cloud Service Resilience using Brownout-Aware Load-Balancing", (submitted to SRDS 2014)
6. L. Tomás, C. Klein, J. Tordsson, F. Hernández-Rodríguez, "The straw that broke the camel's back: safe cloud overbooking with application brownout", (submitted to CAC 2014)

## References

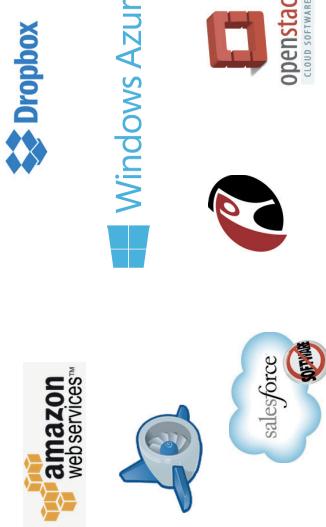
## Results: 2 Applications



37

38

## Cloud Computing



**RT-XEN: REAL-TIME VIRTUALIZATION FOR THE CLOUD****Chenyang Lu, Washington University in St. Louis**

Recent years have witnessed increasing demand of running real-time applications in the cloud. However, existing virtualization platforms cannot provide real-time performance guarantees to virtual machines. This talk will introduce RT-Xen, a real-time virtual machine scheduling framework in the Xen hypervisor. Built based on compositional real-time scheduling theory, RT-Xen realizes a suite of real-time schedulers spanning the design space including global and partitioned multi-core scheduling, fixed and dynamic priority, and different budget management schemes. Our experimental study shows RT-Xen schedulers deliver significant improvement in real-time performance over Xen's existing credit scheduler and explores the tradeoff in real-time scheduler design in virtualized platforms. RT-Xen has been released as open-source software at <https://sites.google.com/site/realtimexen/>. Work is underway to incorporate RT-Xen in the Xen distribution and to integrate RT-Xen with the OpenStack cloud management system. RT-Xen represents a promising step toward real-time cloud computing for latency-sensitive applications.





## RT-Xen: Real-Time Virtualization for the Cloud

Chenyang Lu  
Cyber-Physical Systems Laboratory  
Department of Computer Science and Engineering



### Real-Time Virtualization

- Cars are becoming real-time mini-clouds!
  - Consolidate 100 ECUs → 10 multicore processors.
  - Integrate multiple vendors' systems → common platforms.
  - Must preserve real-time guarantees on a virtualized platform!
- Internet of Things → Cyber-Physical Systems
  - Smart manufacturing, smart transportation, smart grid.
  - Internet-scale sensing and control → real-time cloud computing.
- Cloud gaming
  - Xbox One cloud offloading computation of environmental elements
  - Sony acquired Gaikai, an open cloud gaming platform.

5/9/14

1



- Cars are becoming real-time mini-clouds!

- Consolidate 100 ECUs → 10 multicore processors.
- Integrate multiple vendors' systems → common platforms.
- Must preserve real-time guarantees on a virtualized platform!

- Internet of Things → Cyber-Physical Systems

- Smart manufacturing, smart transportation, smart grid.
- Internet-scale sensing and control → real-time cloud computing.

- Cloud gaming

- Xbox One cloud offloading computation of environmental elements
- Sony acquired Gaikai, an open cloud gaming platform.

1



### Challenges

- Support real-time applications in a virtualized environment
  - Latency **guarantees** to tasks running in virtual machines (VMs).
  - Real-time performance **isolation** between VMs.
- Real-time performance provisioning at different levels
  - Virtualization within a host
  - Communication and I/O
  - Cloud resource management

3



### Virtualization is *not* real-time today

- Existing hypervisors provide no guarantee on latency
  - Xen: credit scheduler, [credit, cap]
  - VMware ESXi: [reservation, share, limitation]
  - Microsoft Hyper-V: [reserve, weight, limit]
- Public clouds lack service level agreement on latency
  - EC2, Compute Engine, Azure, #vCPUs

**Current platforms provision CPU resources, not real-time performance!**

2

5/9/14

## RT-Xen

- Real-time hypervisor based on Xen
  - Real-time VM scheduling
  - Real-time communication

### Build on compositional scheduling theory

- VMs specify resource interfaces
  - Real-time guarantees to tasks in VMs



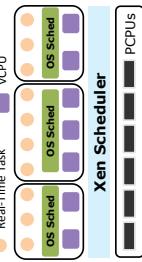
5/9/14 4

## Xen Virtualization Architecture

- Xen type-I, baremetal hypervisor
  - Domain-0: drivers, tool stack to control VMs
  - Guest Domain: para-virtualized or fully virtualized OS.

### Xen scheduler

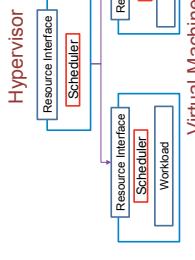
- Guest OS runs on VCPUs.
  - Credit scheduler: round-robin with proportional share.
  - Xen schedules VCPUs on PCPUs.
  - Guest Domain: para-virtualized or fully virtualized OS.



5/9/14 5

## Compositional Scheduling

- Analytical real-time guarantees to tasks running in VMs.
- VM resource interfaces
  - Hides task-specific information
  - Multicore: <period, budget, #CPU>
  - Computed based on compositional scheduling analysis



5/9/14 6

## Real-Time Scheduling Policies

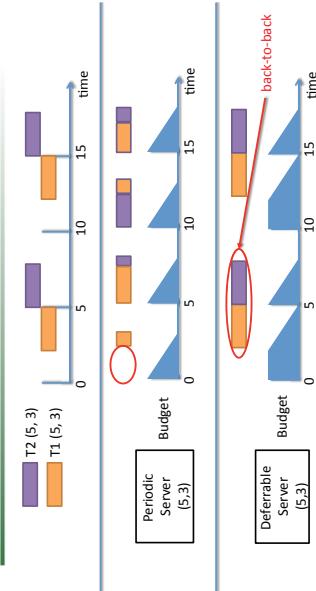
- Priority schemes
  - Static priority: Rate Monotonic
  - Dynamic priority: Earliest Deadline First (EDF)
- Multi-core
  - Global scheduling: allow VCPU migration across cores
  - Partitioned scheduling: bound VCPUs to cores

5/9/14 7



## RT-Xen: Real-Time Scheduling in Xen

### Scheduling VM as "Server"



5/9/14

8

5/9/14

9

### Experimental Setup

- Hardware: Intel i7 processor, 6 cores, 3.33 GHz
  - Allocate 1 vCPU for Domain-0, pinned to PCPU 0
  - All guest VMs use the remaining cores

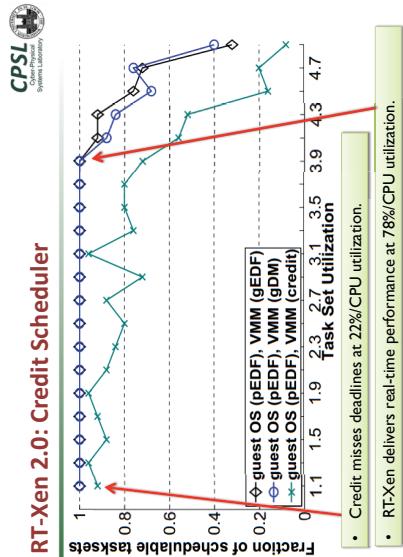
- Software
  - Xen 4.3 patched with RT-Xen
  - Guest OS: Linux patched with LITTMUS

- Workload
  - Period tasks: synthetic ARINC 653 avionics workload (RT-Xen 1.1)
  - Allocate tasks → VMs

5/9/14

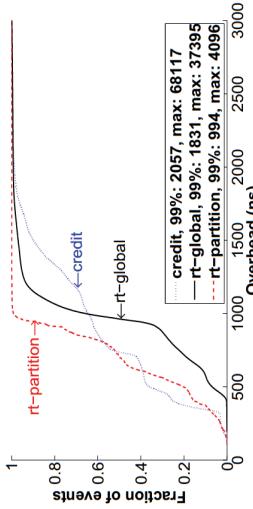
10

### RT-Xen 2.0: Credit Scheduler



11

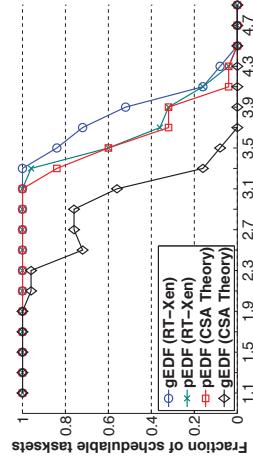
## RT-Xen 2.0: Scheduling Overhead



5/9/14 12

- rt-global has extra overhead due to global lock.
- Credit has poor max overhead due to load balancing.

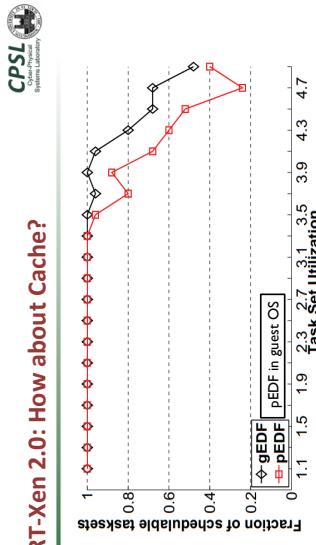
## RT-Xen 2.0: Theory vs. Experiments



5/9/14 13

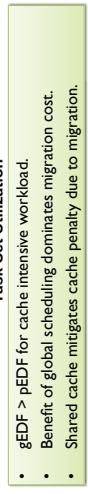
- gEDF > pEDF empirically, thanks to work-conserving global scheduling.
- gEDF < pEDF theoretically due to pessimistic analysis.

## RT-Xen 2.0: Deferrable vs. Periodic



Work-conserving wins empirically!

- Deferrable Server (DS) > Periodic Server.
- gEDF+DS → best real-time performance.



5/9/14 14

- gEDF > pEDF for cache intensive workload.
- Benefit of global scheduling dominates migration cost.
- Shared cache mitigates cache penalty due to migration.

5/9/14 15

## Conclusion

- Diverse applications demand real-time virtualization and cloud.
  - Embedded real-time systems
  - Internet-scale cyber-physical systems
  - Latency-sensitive cloud applications
- RT-Xen provides real-time performance and guarantees
  - Efficient implementation of diverse real-time scheduling policies.
  - Leverage compositional scheduling theory → analytical guarantee.
  - Resource interface → systematic resource allocation for latency bounds.
- On-going
  - Working on RT-Xen patch for Xen core distribution.
  - RT-OpenStack: integration with OpenStack on the way.



## Check out RT-Xen

**RT-Xen**

I'm Feeling Lucky ↗

<https://sites.google.com/site/realtimexen/>

- **RT-Xen 1.0:** S.Xi, Wilson, C. Lu, and C.D. Gill, [RT-Xen Towards Real-Time Hypervisor: Scheduling in Xen](#), ACM International Conferences on Embedded Software (EMSOFT), 2011.
- **RT-Xen 1.1:** Lee, S. Xi, S. Chen, L.T.X. Phan, C. Gill, I. Lee, C. Lu and O. Sokolsky, [Realtime Compositional Scheduling through Virtualization](#), IEEE Real-Time and Embedded Technology and Applications Symposium (RTAS) 2012.
- **RT-Xen 2.0:** S. Xi, C. Lu, C. Gill, M. Xu, L.T.X. Phan, I. Lee, and O. Sokolsky, [Real-Time Multi-Context Virtual Machine Scheduling in Xen](#), Washington University Technical Report, WUCSE-2013-09 (2013)
- **Inter-domain communication:** S. Xi, C. Li, C. Lu, and C. Gill, [Prioritizing Local Inter-Domain Communication in Xen](#), ACM/IEEE International Symposium on Quality of Service (IWQoS), 2013.

17

5/9/14

16

5/9/14

**SERVICE LEVEL AGREEMENT FOR CLOUD COMPUTING:  
TOWARDS A CONTROL-THEORETIC APPROACH**  
**Sara Bouchenak, University of Grenoble**

Cloud Computing is a paradigm for enabling remote, on-demand access to a set of configurable computing resources. This model aims to provide hardware and software services to customers, while minimizing human efforts in terms of service installation, configuration and maintenance, for both cloud provider and cloud customer. A cloud may have the form of an Infrastructure as a Service (IaaS), a Platform as a Service (PaaS) or a Software as a Service (SaaS). However, cloud's ad-hoc management in terms of quality-of-service and service level agreement (SLA) poses significant challenges to the performance, availability, energy consumption and economical costs of the cloud. We believe that a differentiating element between Cloud Computing environments will be the quality-of-service and the service level agreement (SLA) provided by the cloud. In this talk, we will discuss the definition and implementation of a novel cloud model: SLAaaS (SLA aware Service). The SLAaaS model enriches the general paradigm of Cloud Computing, and enables systematic and transparent integration of service levels and SLA to the cloud. SLAaaS is orthogonal to IaaS, PaaS and SaaS clouds and may apply to any of them. Both the cloud provider and cloud customer points of view are taken into account. From cloud provider's point of view, we present autonomic SLA management to handle performance, availability, energy and cost issues in the cloud. An innovative approach combines control theory techniques with distributed algorithms and language support in order to build autonomic elastic clouds. Novel models, control laws, distributed algorithms and languages will be proposed for automated provisioning, configuration and deployment of cloud services to meet SLA requirements, while tackling scalability and dynamics issues. On the other hand from cloud customer's point of view, we discuss SLA governance. It allows cloud customers to be part of the loop and to be automatically notified about the state of the cloud, such as SLA violation and cloud energy consumption. The former provides more transparency about SLA guarantees, and the latter aims to raise customers' awareness about cloud's energy footprint.



## A Control Approach for Performance of Big Data Systems

**Mihaly Berekmeri**, Damian Serrano, Sara Bouchenak,  
Nicolas Marchand, Bogdan Robu

GIPSA - LIG - Grenoble University, France



LCCC 2014, Lund, Sweden

### The structure of the presentation

1. Introduction
  - Big Data MapReduce
  - State of the art
2. Experimental setup
  - Sensors / Actuators
  - MRBS
3. Control
  - Our model
  - Control architecture
  - Control examples
4. Conclusions and Future Work

2

## Big Data, Big Problems



Conclusions

1

## MapReduce

### Problem:

Vast amounts of data generated daily

– Facebook:

- $1.11 \times 10^9$  active users, 50% log in daily
- $3.2 \times 10^6$  likes and comments/day
- > 100 clusters (largest has > 100PB, 200 million files)

– CERN's LHC: Up to 1 PB/s during experiments

### How do we store it? How do we process it?



3

### Used by the biggest companies :

Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...  
**Wide range of applications :**  
log analysis, data mining, web search engines, scientific computing, business intelligence,...

4

LCCC 2014, Lund, Sweden

Ben Chavis - Fotolia

# A Control Approach for Performance of Big Data Systems

Mihaly Berekmeri, Damian Serrano, Sara Bouchenak,  
Nicolas Marchand, Bogdan Robu

GIPSA - LIG - Grenoble University, France



## The structure of the presentation

1. Introduction
  - Big Data MapReduce
  - State of the art
2. Experimental setup
  - Sensors / Actuators
  - MREBS
3. Control
  - Our model
  - Control architecture
  - Control examples
4. Conclusions and Future Work

1

LCCC2014, Lund, Sweden

## Big Data, Big Problems

*Problem:*  
Vast amounts of data generated daily

– Facebook:

- $1.11 \times 10^9$  active users, 50% log in daily
- $3.2 \times 10^6$  likes and comments/day
- > 100 clusters (largest has > 100PB, 200 million files)

– CERN's LHC: Up to 1 PB/s during experiments  
**How do we store it? How do we process it?**



## MapReduce

**Programming model** introduced by J. Dean and S. Ghemawat (Google) in 2004 as a PaaS paradigm  $\rightarrow$  large scale distributed data processing on clusters of commodity computers

**Automatic features:** data partitioning and replication, task scheduling, fault tolerance  
**Used by the biggest companies :**  
Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...  
**Wide range of applications :**  
log analysis, data mining, web search engines, scientific computing, business intelligence,...

Ben Chams - FORGE



LCCC 2014, Lund, Sweden

3

Introduction > Experimental setup > Control > Conclusions

Introduction > Experimental setup > Control > Conclusions

2



4

## A Control Approach for Performance of Big Data Systems

**Mihaly Berekmeri**, Damian Serrano, Sara Bouchenak,  
Nicolas Marchand, Bogdan Robu

GIPSA - LIG - Grenoble University, France



LCCC 2014, Lund, Sweden

### The structure of the presentation

1. Introduction
  - Big Data MapReduce
  - State of the art
2. Experimental setup
  - Sensors / Actuators
  - MRBS
3. Control
  - Our model
  - Control architecture
  - Control examples
4. Conclusions and Future Work

2

## Big Data, Big Problems



### Problem:

Vast amounts of data generated daily

– Facebook:

- $1.11 \times 10^9$  active users, 50% log in daily
- $3.2 \times 10^6$  likes and comments/day
- > 100 clusters (largest has > 100PB, 200 million files)

– CERN's LHC: Up to 1 PB/s during experiments  
**How do we store it? How do we process it?**



Used by the biggest companies :

Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...  
**Wide range of applications :**  
 log analysis, data mining, web search engines, scientific computing, business intelligence,...

LCCC 2014, Lund, Sweden

3

## MapReduce

**Programming model** introduced by J. Dean and S. Ghemawat (Google) in 2004 as a PaaS paradigm  $\rightarrow$  large scale distributed data processing on clusters of commodity computers  
**Automatic features:** data partitioning and replication, task scheduling, fault tolerance

**Used by the biggest companies :**  
 Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...

**Wide range of applications :**  
 log analysis, data mining, web search engines, scientific computing, business intelligence,...

4



2

# A Control Approach for Performance of Big Data Systems

**Mihaly Berekmeri**, Damian Serrano, Sara Bouchenak,  
Nicolas Marchand, Bogdan Robu

GIPSA - LIG - Grenoble University, France



LCCC2014, Lund, Sweden

## Big Data, Big Problems



*Problem:*  
Vast amounts of data generated daily

– Facebook:

- $1.11 \times 10^9$  active users, 50% log in daily
- $3.2 \times 10^6$  likes and comments/day
- > 100 clusters (largest has > 100PB, 200 million files)

– CERN's LHC: Up to 1 PB/s during experiments

**How do we store it? How do we process it?**



Ben Chams - FORCIA

## The structure of the presentation

1. **Introduction**
  - Big Data MapReduce
  - State of the art
2. **Experimental setup**
  - Sensors / Actuators
  - MRBS
3. **Control**
  - Our model
  - Control architecture
  - Control examples
4. **Conclusions and Future Work**

1

## MapReduce

**Programming model** introduced by J. Dean and S. Ghemawat (Google) in 2004 as a PaaS paradigm  $\rightarrow$  large scale distributed data processing on clusters of commodity computers

**Automatic features:** data partitioning and replication, task scheduling, fault tolerance

**Used by the biggest companies :**

Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...  
**Wide range of applications :**  
log analysis, data mining, web search engines, scientific computing, business intelligence,...

2

LCCC 2014, Lund, Sweden



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



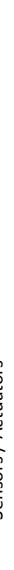
Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



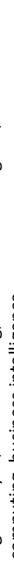
Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions

3

Ben Chams - FORCIA



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions



Introduction > Experimental setup > Control > Conclusions

4

## A Control Approach for Performance of Big Data Systems

**Mihaly Berekmeri**, Damian Serrano, Sara Bouchenak,  
Nicolas Marchand, Bogdan Robu

GIPSA - LIG - Grenoble University, France



LCCC 2014, Lund, Sweden

### The structure of the presentation

1. Introduction
  - Big Data MapReduce
  - State of the art
2. Experimental setup
  - Sensors / Actuators
  - MRBS
3. Control
  - Our model
  - Control architecture
  - Control examples
4. Conclusions and Future Work

2

## Big Data, Big Problems



### Big Data, Big Problems

*Problem:*  
Vast amounts of data generated daily

- Facebook:
  - $1.11 \times 10^9$  active users, 50% log in daily
  - $3.2 \times 10^6$  likes and comments/day
  - > 100 clusters (largest has > 100PB, 200 million files)
- CERN's LHC: Up to 1 PB/s during experiments

### How do we store it? How do we process it?



Used by the biggest companies :

Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...  
Wide range of applications :

log analysis, data mining, web search engines, scientific computing, business intelligence,...

LCCC 2014, Lund, Sweden

4

## MapReduce

**Programming model** introduced by J. Dean and S. Ghemawat (Google) in 2004 as a PaaS paradigm  $\rightarrow$  large scale distributed data processing on clusters of commodity computers  
**Automatic features:** data partitioning and replication, task scheduling, fault tolerance

**Used by the biggest companies :**

Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...  
Wide range of applications :

log analysis, data mining, web search engines, scientific computing, business intelligence,...

Ben Chavis - Fotolia

3

# A Control Approach for Performance of Big Data Systems

**Mihaly Berekmeri**, Damian Serrano, Sara Bouchemak,  
Nicolas Marchand, Bogdan Robu

GIPSA - LIG - Grenoble University, France



## Big Data, Big Problems

*Problem:*  
Vast amounts of data generated daily

– Facebook:

- $1.11 \times 10^9$  active users, 50% log in daily
- $3.2 \times 10^6$  likes and comments/day
- > 100 clusters (largest has > 100PB, 200 million files)

– CERN's LHC: Up to 1 PB/s during experiments  
**How do we store it? How do we process it?**



## The structure of the presentation

- Introduction**
  - Big Data MapReduce
  - State of the art
- Experimental setup**
  - Sensors / Actuators
  - MRBS
- Control**
  - Our model
  - Control architecture
  - Control examples
- Conclusions and Future Work**

## MapReduce

**Programming model** introduced by J. Dean and S. Ghemawat (Google) in 2004 as a PaaS paradigm  $\rightarrow$  large scale distributed data processing on clusters of commodity computers

**Automatic features:** data partitioning and replication, task scheduling, fault tolerance

**Used by the biggest companies :**

Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...

### Wide range of applications :

log analysis, data mining, web search engines, scientific computing, business intelligence,...

LCCC 2014, Lund, Sweden

Introduction > Experimental setup > Control > Conclusions

## A Control Approach for Performance of Big Data Systems

**Mihaly Berekmeri**, Damian Serrano, Sara Bouchenak,  
Nicolas Marchand, Bogdan Robu

GIPSA - LIG - Grenoble University, France



ICCC'2014, Lund, Sweden

### The structure of the presentation

1. Introduction
  - Big Data MapReduce
  - State of the art
2. Experimental setup
  - Sensors / Actuators
  - MRBS
3. Control
  - Our model
  - Control architecture
  - Control examples
4. Conclusions and Future Work

2

## Big Data, Big Problems

*Problem:*  
Vast amounts of data generated daily

- Facebook:
  - $1.11 \times 10^9$  active users, 50% log in daily
  - $3.2 \times 10^6$  likes and comments/day
  - > 100 clusters (largest has > 100PB, 200 million files)
- CERN's LHC: Up to 1 PB/s during experiments

### How do we store it? How do we process it?



Ben Chems - Fotolia

- Used by the biggest companies :**  
Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...  
**Wide range of applications :**  
log analysis, data mining, web search engines, scientific computing, business intelligence,...

4

## MapReduce

**Programming model** introduced by J. Dean and S. Ghemawat (Google) in 2004 as a PaaS paradigm → large scale distributed data processing on clusters of commodity computers  
**Automatic features:** data partitioning and replication, task scheduling, fault tolerance

**Used by the biggest companies :**  
Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...

- Used by the biggest companies :**  
Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...  
**Wide range of applications :**  
log analysis, data mining, web search engines, scientific computing, business intelligence,...

ICCC'2014, Lund, Sweden

5

## A Control Approach for Performance of Big Data Systems

**Mihaly Berekmeri**, Damian Serrano, Sara Bouchenak,  
Nicolas Marchand, Bogdan Robu

GIPSA - LIG - Grenoble University, France



### Big Data, Big Problems



**Problem:**  
Vast amounts of data generated daily

– Facebook:

- $1.11 \times 10^9$  active users, 50% log in daily
- $3.2 \times 10^8$  likes and comments/day
- > 100 clusters (largest has > 100PB, 200 million files)

– CERN's LHC: Up to 1 PB/s during experiments

**How do we store it? How do we process it?**



## The structure of the presentation

1. **Introduction**
  - Big Data MapReduce
  - State of the art
2. **Experimental setup**
  - Sensors / Actuators
  - MRBS
3. **Control**
  - Our model
  - Control architecture
  - Control examples
4. **Conclusions and Future Work**

LCCC2014, Lund, Sweden

1

### MapReduce

**Programming model** introduced by J. Dean and S. Ghemawat (Google) in 2004 as a PaaS paradigm  $\rightarrow$  large scale distributed data processing on clusters of commodity computers  
**Automatic features:** data partitioning and replication, task scheduling, fault tolerance

**Used by the biggest companies :**  
Amazon, eBay, Facebook, LinkedIn, Twitter, Yahoo, Microsoft...

**Wide range of applications :**  
log analysis, data mining, web search engines, scientific computing, business intelligence,...

LCCC 2014, Lund, Sweden

3

Ben Chams - Fossile

4

**PERFORMANCE-ENERGY TRADE-OFF IN MULTI-SERVER  
QUEUEING SYSTEMS WITH SETUP DELAY**  
**Samuli Aalto, Aalto University**

In this talk we review some recent results related to the performance-energy trade-off in multi-server queueing systems, where the servers have multiple energy states. In addition to the normal BUSY and IDLE states, a server can be switched OFF to save energy. However, switching the server again on results in a SETUP delay which consumes additional energy and deteriorates performance. For a single server system, we consider optimal strategies to switch the server off and on. Multi-server systems may have a central queue, or they may consist of parallel servers with their own queues. Switching servers off and on in a reasonable way is the main objective. In a system with parallel servers, one should also consider the dispatching (a.k.a. task assignment) problem: an arriving job is to be routed to one of the parallel servers. Here we present a size- and energy-aware MDP approach to solve the problem.

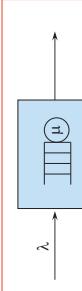


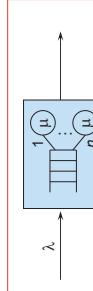
# Performance-Energy Trade-off in Multi-Server Queueing Systems with Setup Delay

Samuli Aalto  
 Aalto University, Finland

LCCC Cloud Computing Workshop  
 7–9 May 2014  
 Lund, Sweden

## Queueing models

- Single-server queue (M/G/1)
- 

- Multi-server queue (M/M/n)
- 

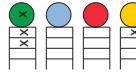
## Co-operation with

Esa Hyttiä, Pasi Lasila, Misikir Gebrehiwot  
 (Aalto University)

Rhonda Righter  
 (UC Berkeley)

## Performance-energy trade-off

BUSY / IDLE / OFF / SETUP?



- Energy saved by switching the server off when idle
- However, performance impaired, if switching the server back on takes time (**setup delay**)

## Cost model

- Performance:
  - Energy:  
 $E[E] = \text{mean energy per job}$   
 (in joules)
- $E[X] = \text{mean number of jobs}$   
 $= \lambda \cdot E[T]$
- Definition:  
**delay = response time**

- Energy-time-Weighted-Sum (ERWS):  
 $E[T] + E[E]\beta$
- $E[P] = \text{mean power consumed}$   
 $= \lambda \cdot E[E]$
- Power consumption levels:  
 $0 = P_{\text{off}} < P_{\text{idle}} \leq P_{\text{setup}} = P_{\text{busy}}$
- **Theorem:**  
**delay = response time**

- Energy-Response-time-Weighted-Sum (ERWS):  
 $w_1 \cdot E[T]^{\alpha_1} \cdot E[E]^{\alpha_1} + w_2 \cdot E[T]^{\alpha_2} \cdot E[E]^{\alpha_2}$
- e.g. Wieman & al. (2009)
- Energy-Response-time-Product (ERP):  
 $E[T] \cdot E[E]$
- e.g. Gandhi & al. (2010b)

## Objective function

- General form:
- Energy-Response-time-Weighted-Sum (ERWS):  
 $E[T] + E[E]\beta$
- e.g. Maccio & Down (2013)
- Energy-Response-time-Product (ERP):  
 $E[T] \cdot E[E]$
- e.g. Maccio & Down (2013)



Introduction

5

Introduction

6



## Part I Single-server queue with setup delays

### Optimal switching on/off policy

Maccio & Down (2013)  
 Auto University  
School of Electrical  
Engineering

- Policies:
  - M/G/1-FIFO
  - Setup delay  $D$  generally distributed with mean  $1/\gamma$
- Control parameters:
  - Delayed switch-off for an exponential time with mean  $1/\alpha$
  - Server switched on after  $K$  new job arrivals
- Theorem:
  - For ERWS objective function optimal policy is either NEVEROFF or INSTANTOFF
- Objective function: Gen. form  
 $w_1 E[T]^{\alpha_1} E[E]^{\alpha_1} + w_2 E[T]^{\alpha_2} E[E]^{\alpha_2}$



Part I  
Single-server queue

7

Part I  
Single-server queue

8

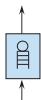


Part I  
Single-server queue

## Optimal switching on/off policy

Gebrehiwot & al. (2014)

- Policies:
  - NEVEROFF:  $\alpha = 0$
  - DELAYEDOFF:  $0 < \alpha < \infty$
  - INSTANTOFF:  $\alpha = \infty$
- Control parameters:
  - Delayed switch-off for a gen. distributed time with mean  $1/\alpha$
  - Server switched on after  $k$  new job arrivals
- Objective function: Gen. form
 
$$w_1 E[T]^t + w_2 E[T]^2 E[E]^{e_2}$$
- Theorem:
  - For gen. objective function optimal policy is either NEVEROFF or INSTANTOFF



## Part II Multi-server queue with setup delays

M/G/1-FIFO

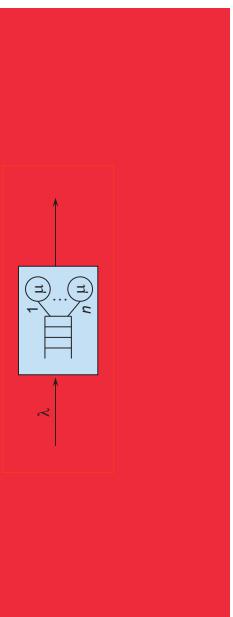
- Setup delay  $D$  generally distributed with mean  $1/\gamma$

Control parameters:

- Delayed switch-off for a gen. distributed time with mean  $1/\alpha$
- Server switched on after  $k$  new job arrivals

• Objective function: Gen. form

- NEVEROFF is better if  $P_{idle}$  is sufficiently small compared to  $P_{setup}$



**A**

Part I  
Single-server queue

**A** Auto University  
School of Electrical  
Engineering



## Analysis of server farms with setup delays

Gandhi & al. (2010a)

- M/M/m
  - Setup delay  $D$  exponentially distributed
- Objective function:
  - Separately  $E[T]$  and  $E[P]$

- Policies:
  - ON/IDLE = NEVEROFF
  - ON/OFF = INSTANTOFF
  - ON/OFF/STAG = INSTANTOFF with staggered bootup

- Mixed policy:
  - ON/IDLE(t)
  - switching idle server off only if nr of busy and idle servers > t
- Conclusions:
  - Under high loads, turning servers off can result in higher power consumption and far higher response times."
  - "As the size of the server farm is increased, the advantages of turning servers off increase."

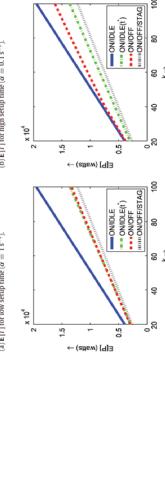
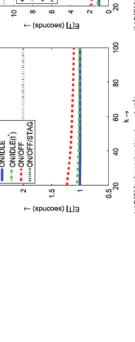
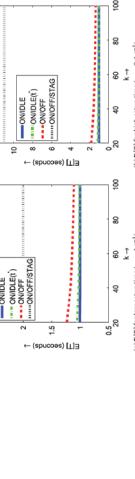


**A** Auto University  
School of Electrical  
Engineering



## Analysis of server farms with setup delays

Gandhi & al. (2010a)



Part II  
Multi-server queue

**A** Auto University  
School of Electrical  
Engineering

**A** Auto University  
School of Electrical  
Engineering

## Optimization of server farms with setup delay

Gandhi &amp; al. (2010)b

- M/M/n
  - Setup delay deterministic
  - Additional sleep states  $S$  with  $0 = P_{\text{off}} < P_{\text{sleep}} < P_{\text{idle}}$
- Policies:
  - NEVEROFF
  - INSTANTOFF
  - SLEEP(S)
  - Probabilistic and other
- Theorem:
 

For  $n = 1$ , optimal static control is either NEVEROFF, INSTANTOFF or SLEEP( $S$ )
- Robust policy:
  - DELAYEDOFF with MRB (Most Recent Busy)

Aalto University  
School of Electrical  
Engineering  
Part II  
Multi-server queue  
13

Aalto University  
School of Electrical  
Engineering  
Part II  
Multi-server queue  
14

## Optimization of server farms with setup delay

Gandhi &amp; al. (2010)b

- Policies:
  - NEVEROFF
  - INSTANTOFF
  - SLEEP(S)
  - Probabilistic and other
- Theorem:
 

For  $n = 1$ , optimal static control is either NEVEROFF, INSTANTOFF or SLEEP( $S$ )
- Robust policy:
  - DELAYEDOFF with MRB (Most Recent Busy)

Aalto University  
School of Electrical  
Engineering  
Part II  
Multi-server queue  
13

Aalto University  
School of Electrical  
Engineering  
Part II  
Multi-server queue  
14

## Optimization of server farms with setup delay

Gandhi &amp; al. (2010)b

- Policies:
  - NEVEROFF
  - INSTANTOFF
  - SLEEP(S)
  - Probabilistic and other
- Theorem:
 

For  $n = 1$ , optimal static control is either NEVEROFF, INSTANTOFF or SLEEP( $S$ )
- Robust policy:
  - DELAYEDOFF with MRB (Most Recent Busy)

Aalto University  
School of Electrical  
Engineering  
Part II  
Multi-server queue  
13

Aalto University  
School of Electrical  
Engineering  
Part II  
Multi-server queue  
14

## Optimization of server farms with setup delay

Gandhi &amp; al. (2010)b

- Policies:
  - NEVEROFF
  - INSTANTOFF
  - SLEEP(S)
  - Probabilistic and other
- Theorem:
 

For  $n = 1$ , optimal static control is either NEVEROFF, INSTANTOFF or SLEEP( $S$ )
- Robust policy:
  - DELAYEDOFF with MRB (Most Recent Busy)

Aalto University  
School of Electrical  
Engineering  
Part II  
Multi-server queue  
13

Aalto University  
School of Electrical  
Engineering  
Part II  
Multi-server queue  
14

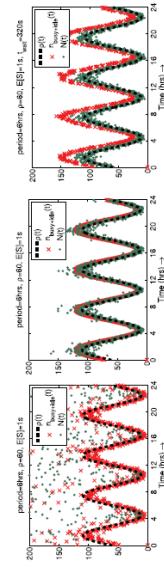
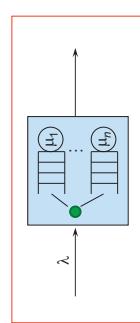
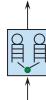


Fig. 4. Dynamic capacity provisioning capabilities of INSTANTOFF, LONGHEAD and NEVEROFF. The dashed line denotes the load at time  $t = t_{\text{avg}} + \Delta t$ , the crosses denotes the number of servers at time  $t = t_{\text{avg}} + \Delta t$ , and the dots represent the number of pos in the system at time  $t = t_{\text{avg}}$ .



## Dispatching problem

- Dispatching
  - = Task assignment = Routing
  - Random job arrivals with random service requirements
  - Dispatching decision made upon the arrival

### Our setting: M/G/1.

- Poisson arrivals
- Generally distributed job sizes
- heterogeneous servers with FIFO queuing discipline (NEVEROFF or INSTANTOFF)

## Part III Parallel queues with setup delays

Part III  
Parallel queues

15

Aalto University  
School of Electrical  
Engineering  
Part III  
Parallel queues

Aalto University  
School of Electrical  
Engineering  
Part III  
Parallel queues

Part III  
Parallel queues

16

Part III  
Parallel queues

17

Aalto University  
School of Electrical  
Engineering  
Part III  
Parallel queues

## Static dispatching policies



### MDP approach

- RND = Bernoulli splitting
  - choose the queue pure randomly
  - **no size nor state information needed**
  - Harchol-Balter et al. (1999)
- SITA = Size Interval Task Assignment
  - choose the queue with similar jobs
  - **based on the size of the arriving job, but no state information needed**

- Any **static policy** (RND, SITA) results in parallel M/G/1 queues
- Fix the static policy and determine **relative values** for all these parallel M/G/1 queues
- Dispatch the arriving job to the queue that **minimizes** the mean additional costs
- As the result, you get a better **dynamic dispatching policy**
- This is called **First Policy Iteration (FPI)** in the MDP theory
- Applicable for the **ERWS** objective function

## A

**A** Auto University  
School of Electrical Engineering

Part III  
Parallel queues

17

## Relative values

- **Definition:** For a fixed policy resulting in a stable system, the **relative value**  $v(x) - v(0)$  gives the expected difference in the infinite horizon cumulative costs between
  - the system initially in state  $x$ , and
  - the system initially in equilibrium

- Mean values:
 
$$E[T] = E[S] + \frac{\lambda E[S^2]}{2(1-\rho)}$$

$$E[P] = (1-\rho)P_{\text{idle}} + \rho \cdot P_{\text{busy}}$$
- **Result:** Size-aware relative values
 
$$v_T(u) - v_T(0) = \frac{\lambda u^2}{2(1-\rho)}$$

$$v_P(u) - v_P(0) = u \cdot (P_{\text{busy}} - P_{\text{idle}})$$

**S** Hyttia et al. (2012)



## Size-aware M/G/1 queue without setup delays

**A** Auto University  
School of Electrical Engineering

18

Part III  
Parallel queues



## Size-aware M/G/1 queue with setup delays

**A** Auto University  
School of Electrical Engineering

19

Part III  
Parallel queues

Part III  
Parallel queues

**A** Auto University  
School of Electrical Engineering

20

## Size-aware M/G/1 queue with setup delays

**Hyytiä et al. (2014a)**

- State description:

$$u = \Delta_0 + \Delta_1 + \dots + \Delta_n$$

–  $\Delta_i$  = remaining service time of job  $i$

–  $\Delta_0$  = remaining setup delay

–  $u$  = virtual backlog

- Assume:  
Deterministic setup delay  $d$  and

$$P_{\text{setup}} = P_{\text{busy}}$$

- Mean values:

$$E[T] = E[S] + \frac{2E[S^2]}{2(1-\rho)} + \frac{d(2+\lambda d)}{2(1-\rho)}$$

$$E[P] = \frac{\rho + \lambda d}{1 + \lambda d} \cdot P_{\text{busy}}$$

- Result:

Size-aware relative values

$$\nu_T(u) - \nu_T(0) = \frac{\lambda d u^2}{2(1-\rho)} - \frac{\lambda u d(2+\lambda d)}{2(1-\rho)(1+\lambda d)}$$

$$\nu_P(u) - \nu_P(0) = \frac{u}{1+\lambda d} \cdot P_{\text{busy}}$$



Aalto University  
School of Electrical  
Engineering

Part III  
Part and queues

22

## Numerical results

**Hyytiä et al. (2014a)**

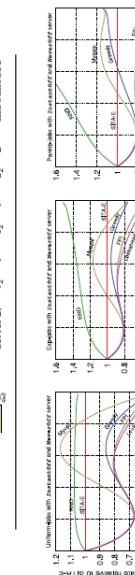


**Table 2**  
Two-server system.

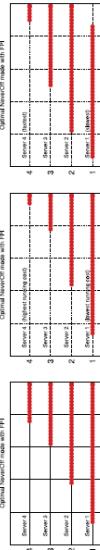
	Server 1: $v_1 = 1$	Server 2: $v_2 = 1$	NeverOff	InstantOff
$d_1 = 1$	$d_2 = 1$	$d_1 = 2$	$d_2 = 2$	$d_1 = 2$

**Table 3**  
Four-server systems.

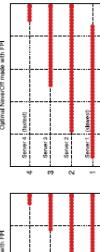
Parameter	(a) Identical	(b) Bilinear	(c) Squared
Service rates	$v_1, \dots, v_4$	$1, 1, 1, 1$	$1, 1, 1, 1$
Running costs	$c_1, \dots, c_4$	$1, 1, 1, 1$	$1, 2, 3, 4$
Switching delay	$d_1, \dots, d_4$	$1, 1, 1, 1$	$1, 2, 9, 16$
			$1, 1, 1, 1$



(a) Uniform.  
(b) Exponential.  
(c) Truncated Pareto.



(b) Exponential.  
(c) Truncated Pareto.



(c) Truncated Pareto.

## FPI policy

**Hyytiä et al. (2012, 2014a)**

- For NEVEROFF servers:

Dispatch the job with service time  $x$  to queue  $i$  minimizing the mean additional costs:

$$a_T(u, x, i) = u + x + d_i \cdot 1(u=0) + \\ \nu_T(u+x, i) - \nu_T(u, i) \\ a_P(u, x, i) = \\ \nu_P(u+x, i) - \nu_P(u, i)$$



Part III  
Part and queues

22

## Numerical results

**Hyytiä et al. (2014a)**



Part III  
Part and queues

22



Aalto University  
School of Electrical  
Engineering

Part III  
Part and queues

24



Aalto University  
School of Electrical  
Engineering

Part III  
Part and queues

## Numerical results

Hyttia et al. (2014a)



## Other queueing disciplines

Hyttia et al. (2014b)



**Table 3**  
Four server systems.

	Parameter	(a) Identical	(b) Linear $\epsilon$	(c) Squared $\epsilon$
Service rates	$v_1, \dots, v_4$	1, 1, 1, 1	1, 1, 1, 1	1, 1, 1, 1
Running costs	$\epsilon_1, \dots, \epsilon_4$	1, 1, 1, 1	1, 2, 3, 4	1, 2, 9, 16
Switching delay	$d_1, \dots, d_4$	1, 1, 1, 1	1, 1, 1, 1	1, 1, 1, 1

(a) Identical servers.

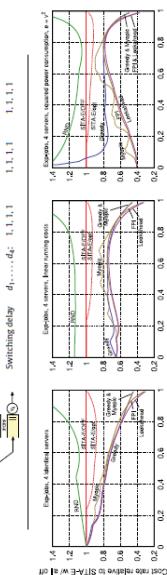
(b) Linear running cost.

(c) Squared running cost.

Aalto University  
School of Electrical  
Engineering

Part III  
Parallel queues

Aalto University  
School of Electrical  
Engineering



Aalto University  
School of Electrical  
Engineering

Part III  
Parallel queues

Aalto University  
School of Electrical  
Engineering

- But it is another story ...

- LIFO in the M/G/ $r$  setting with setup delays
- PS in the M/D/ $r$  setting with setup delays

**Table 3**  
Four server systems.

- LIFO in the M/G/ $r$ , setting with setup delays
- PS in the M/D/ $r$ , setting with setup delays

## References

- Harchol-Balter, Crovella & Murta (1999)  
On Choosing a Task Assignment Policy for a Distributed Server System, *JPDC*
- Weiman, Andrew & Tang (2009)  
Power-aware speed scaling in processor sharing systems, in *IEEE INFOCOM*
- Gandhi, Harchol-Balter & Adam (2010a)  
Server farms with setup costs, *PEVA*
- Gandhi, Gupta, Harchol-Balter & Kozachik (2010b)  
Optimality analysis of energy-performance trade-off for server farm management, *PEVA*
- Hyttia, Penttinen & Aalto (2012)  
Size- and state-aware dispatching problem with queue-specific job sizes, *EJOR*
- Macco & Down (2013)  
Optimal policies for energy-aware servers, in *MASCOTS*
- Hyttia, Raithe & Aalto (2014a)  
Task assignment in a heterogeneous server farm with switching delays and general energy-aware cost structure, to appear in *PEVA*
- Hyttia, Raithe & Aalto (2014b)  
Energy-aware job assignment in server farms with setup delays under LCFS and PS, accepted to *TCC*
- Gebrehiwot, Lassila & Aalto (2014)  
Energy-aware queueing models and controls for server farms, ongoing work

Aalto University  
School of Electrical  
Engineering

References

Aalto University  
School of Electrical  
Engineering

Aalto University  
School of Electrical  
Engineering

Aalto University  
School of Electrical  
Engineering

## The End

**LCCC ACTIVITIES IN CLOUD CONTROL****Maria Kihl, Lund University**

Cloud Control is a large framework project with researchers from both LCCC and the Cloud research group at Umeå University. The main objective with the project is to solve a range of cloud management problems with a control theoretic approach. In this presentation,

I will give an overview of some of the current work in Cloud Control within LCCC.



**LCCC activities in Cloud Control**

MARIA KIHL

## Maria Kihl

- PhD in Teletraffic Systems, 1999
- Associate Professor at Dept. of Electrical and Information Technology since 2004.
- Internet related research projects, often with industry collaboration.
- Joint work with Dept. of Automatic Control since 2002.

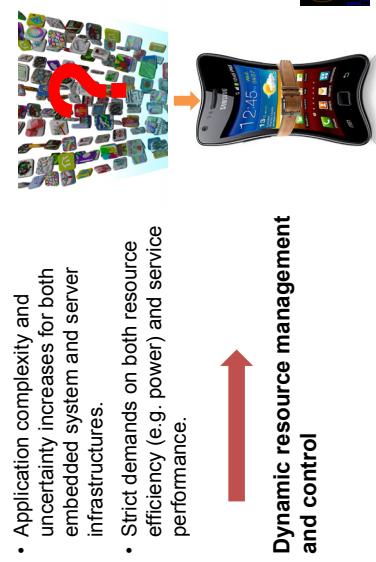


**Main research interests:**  
performance modelling, analysis  
and control of **internetworked**  
systems.

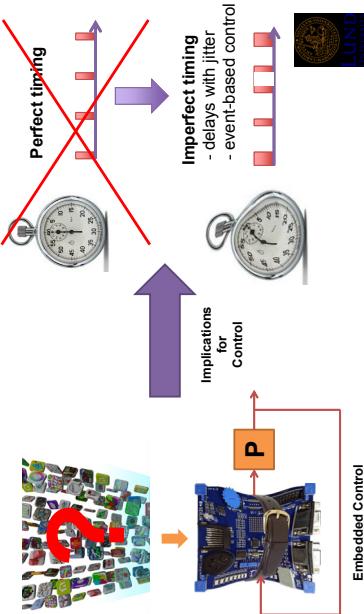
## Implementation in Networked and Embedded Systems

- Maria Kihl  
Karl Erik Arzen  
Anders Robertsson  
Bo Bernhardsson  
Johan Eker (LU, Ericsson AB)  
Martina Maggio  
Anton Cervin  
Alessandro Papadopoulos  
Manfred Dillkrautz (PhD student)  
Jonas Dürango (PhD student)  
William Tärneberg (PhD student)  
Yang Xu (PhD student)  
Mikael Lindberg (PhD stud)  
Payam Amani (PhD stud)

## Background



## Networked Embedded Control



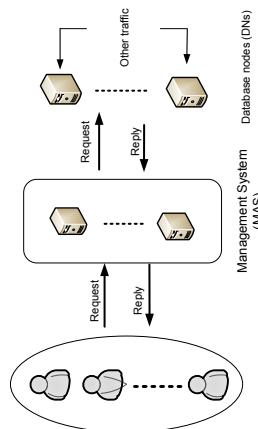
## Resource optimization and control of server systems with latency constraints



- Performance models
- Admission control
- Prediction based capacity optimization

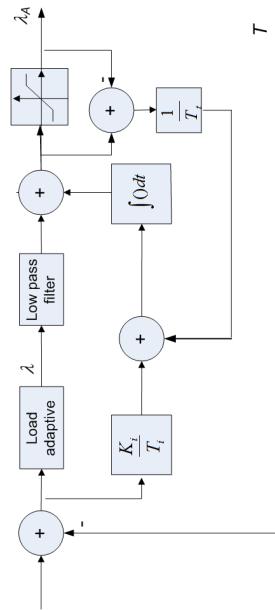


## Example: Mobile Service Support Systems developed by Ericsson AB

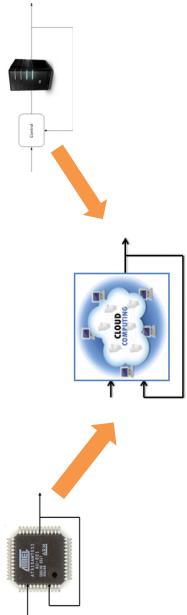


Load control of database nodes with unknown high priority background traffic.

## Example: Load adaptive controller for mobile service support systems

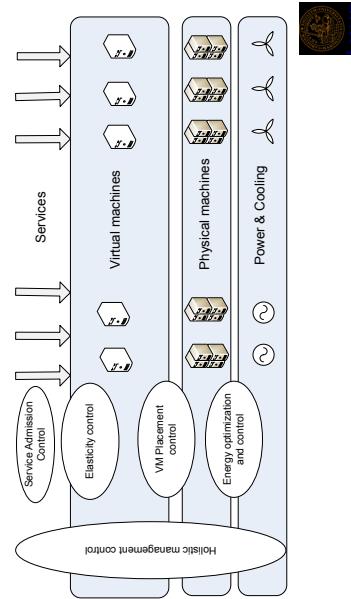


## Cloud Control



- Joint work using our competences in both embedded control and server systems
- A control theoretic approach to a range of cloud management problems.

## Challenges for Cloud Control



## Some ongoing work

- Brownout
  - Martina Maggio, Alessandro Papadopoulos, Jonas Dürango, Manfred Delkraatz
- VM startup time compensation
  - Manfred Delkraatz
- Omnipresent clouds
  - William Tämeberg
- Modeling and autoscaling
  - Jonas Dürango

## Brownout

(Martina Maggio, Alessandro Papadopoulos,  
Jonas Dürango, Manfred Delkraatz)

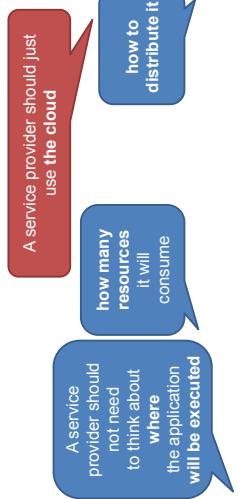
- Cloud applications need to cope with many unexpected events:
  - Flash crowds
  - Hardware failures
  - Unexpected performance degradations

Objective: Applications that withstand variations  
and have similar behaviors in similar conditions



## Motivation: The need for predictability

- Tools like **cloud operating systems** and **virtualization** greatly simplify the development, management and deployment of software applications



## Application performance challenge

However, in reality, another type of problems arise when an application is deployed in the cloud...

Due to a varying amount of physical resources assigned to an application



Use resource allocation mechanisms inspired from embedded systems .

## VM startup time compensation (Manfred Dellekrantz)

Virtual machines take time to start up.  
Controller saying, “Give me  $m$  VMs!” will have to wait for control signal to have effect.

## VM Startup Time Compensation

- VM startup time (“dead time” in control lingo) is a major challenge.
  - Controller reacts several times before its first action has an effect
  - Dead time-unaware controller gives overshoot (unnecessary costs)



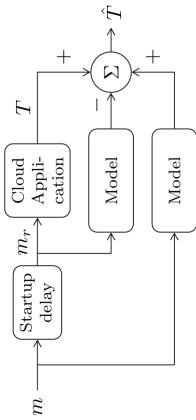
Dead time



## Dead time compensation

- Dead time compensation [Smith, 1957] calculates what the output would have been without dead time.
- Allows you to do faster control, still maintaining stability and avoiding oscillations.
- Requires a model!

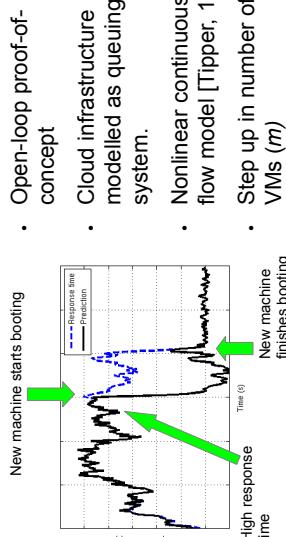
## Dead time compensator



- $m$  – wanted # machines,
- $m_r$  – actual # machines
- $T$  – response time
- $\hat{T}$  – compensated response time



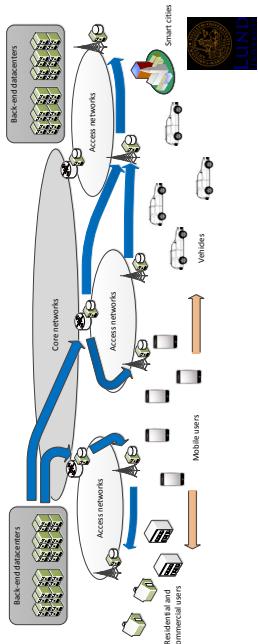
## Proof-of-concept



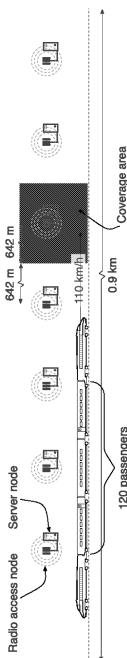
## Omnipresent clouds

(William Tärneberg)

In the omnipresent cloud paradigm, cloud compute resources are migrated or located to the capillaries of the mobile networks.

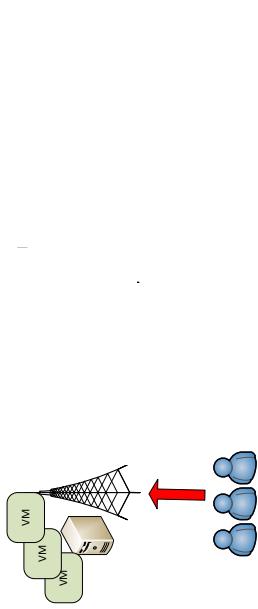


## Scenario



- Cellular network corresponding to 4G LTE.
- First mobility model corresponding to a train.
  - Easily extended to a highway with cars.
- Users request a web-like cloud service according to a stochastic process.
- VMs can be migrated between cells with a certain delay.

## Cloud capacity moves with users

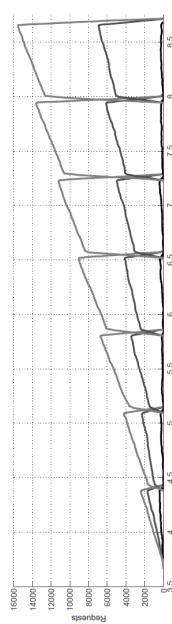


## Cloud control challenges

- Lack of good system models
  - Architecture models
  - Service models
  - Mobility models
  - Workload models
- Workload and Capacity demands in both time and space.
- Auto-scaling: When and Where?
- VM migration: Needs to be performed fast enough dependent on the mobility of the users.
- Heterogeneous nodes must be taken into account.



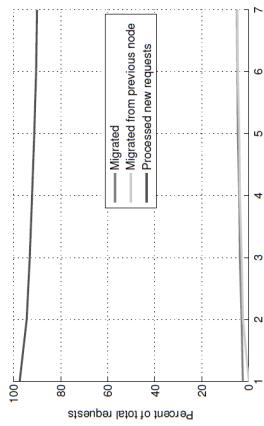
## Load displacement example



- The workload will move between cloud servers, dependent on the mobility of the users.
- An underprovisioned system will spread in both time and space.



## Performance during normal load



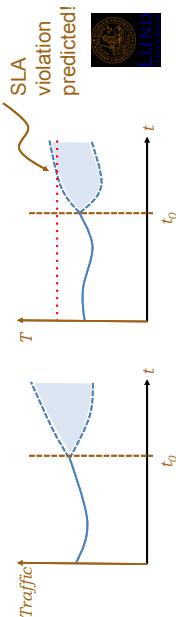
During normal load, the system migrates VMs according to the mobility and handles requests with low response times.



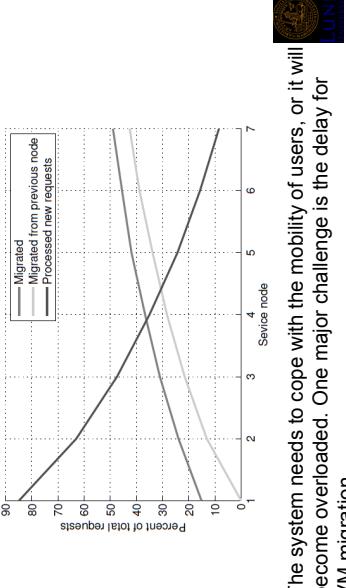
### Modeling and autoscaling (Jonas Dürango)

Computer system models suitable for control are usually tricky to find.

Understanding of dominant system dynamics is vital for decision making (Scale up? Down?)



## Performance during overload



The system needs to cope with the mobility of users, or it will become overloaded. One major challenge is the delay for VM migration.

### Modeling and autoscaling

Use system identification to find models for different purposes:

- Simulation (model-based feedback control)
  - Better understanding of system dynamics -> allows better tuning of controllers.
- Prediction (proactive control)
  - Possibility to anticipate potential SLA violations



## Summary

---

- Within LCCC, we have established a cross-disciplinary research environment with competence in control theory, real-time systems, and internetworked systems.
- Ongoing work on different aspects of Cloud Control.

## Bring theory to practice



## CAPACITY MANAGEMENT IN IAAS CLOUD

**David Breitgand, IBM Research Haifa**

One of the promises of elastic cloud computing is relieving its customers from capacity planning by adding just the right amount of resources just in time when elastic applications need them. While realizing this vision might indeed exempt the customer from the complex and effort consuming capacity management task, the cloud provider still needs to execute on capacity planning to strike the right balance between SLA commitments and cost efficiency. The cost efficiency is intimately related to statistical multiplexing of workloads in the cloud, allowing over-committing cloud resources. Naturally, over-committing implies risk of resource congestion.

Therefore, there is a tradeoff between improving resource utilization by increasing an over-commit ratio and exposing the infrastructure provider and customers to the risk of resource congestion. In this talk I am going to explore a number of approaches to managing this trade-off.



## Focus period and Workshop in Cloud Control

Lund University, Sweden, May 6–10, 2014

### Capacity Management in IaaS Cloud

David Breitgand, IBM Haifa Research Lab



Based on joint works with many collaborators

© 2014 IBM Corporation



### Speaker Background

- Distributed Computing, Hebrew University, Israel
- Networking, Technion, Israel
- Applied Mathematics, Polytechnic University, Novosibirsk, Russia
- Technical Lead of Cloud Operating System Technology Group @IBM Research – Haifa
- Technical Lead of Software Defined Manufacturing Group @IBM Research – Haifa
- 10 years with IBM research, working on algorithms, performance analysis, load balancing, root cause analysis, data center optimization, SLA management, SAN performance management, Cloud computing, VM migration optimization, capacity management.

- **Overarching research theme:** studying, modeling, and optimizing tradeoffs between the costs of management and benefits accrued

2

© 2014 IBM Corporation



### Speaker Background

- Distributed Computing, Hebrew University, Israel
- Networking, Technion, Israel
- Applied Mathematics, Polytechnic University, Novosibirsk, Russia
- Technical Lead of Cloud Operating System Technology Group @IBM Research – Haifa
- Technical Lead of Software Defined Manufacturing Group @IBM Research – Haifa
- 10 years with IBM research, working on algorithms, performance analysis, load balancing, root cause analysis, data center optimization, SLA management, SAN performance management, Cloud computing, VM migration optimization, capacity management.

- **Overarching research theme:** studying, modeling, and optimizing tradeoffs between the costs of management and benefits accrued

2

© 2014 IBM Corporation

### Outline

- Introduction
- Problems
- Results
- Future: where does this become really hard and should we go there?
- Q&A

Introduction

Cloud success == three things:

© 2015 IBM Corporation



Business Model



Simplicity



Agility

© 2015 IBM Corporation

4

© 2015 IBM Corporation

3

introduction

---

## Do I need [long term] Capacity Management in an IaaS Cloud?



Capacity [long term] Management on a IaaS Cloud?

introduction

---

## IBM

introduction

---

## IBM

---

## Capacity [long term] Management on a IaaS Cloud?

© 2010 IBM Corporation

6

introduction

---

## IBM

---

## Capacity [long term] "capacity management" means?

© 2010 IBM Corporation

7

# YES

If you are a **provider** of a public cloud

© 2010 IBM Corporation

8

- IaaS commoditize → IaaS providers operate under perfect competition
- Provider needs to be competitive → Pressure to lower prices
- **Pressure to lower prices** → **Pressure "to do more with less"**
- Capacity management is required to achieve that

▪ The most obvious link to this workshop: cost efficient elasticity

© 2010 IBM Corporation

## How to do more with less?

- By increasing *over-commit ratio* to improve *cost efficiency*,



- But if I *over-commit* too much, I will *degrade cost-efficiency, won't I?*

9

## What does "cost-efficiency" means?

- Cost efficiency [first usage in 1970]
  - "Cost effective describes something that is of a good value, where the benefits and usage are worth at least what is paid for them"

- Cost efficiency is a relative concept describing relationship between **at least** two options

- Example 1:

- Use of PSTN line versus VoIP for long distance call

- Example 2:

- Static resource provisioning versus on-demand provisioning from shared multiplexed pool

## The Gist of "capacity management" for provider

- Find minimal physical resources configuration to provide service to the customer, where the value of using the service worthy at least of what the customer pays for it

- Service in IaaS:

- Provide VM, storage, networks (i.e., resource collections)

- Trade-off:

- Risk of resource congestion vs. cost of using successfully acquired resources

11

---

Introduction

If you expected an answer...  




[Second Price Auctions?](#)

- Vickrey-Clark-Groves
- Heterogeneous goods
- Dominant strategy is to report the true value for the goods
- Beautiful game-theoretic mechanism
- Rarely used in practice
- Complex implementation
- Close to zero seller revenue

13

© 2010 IBM Corporation

Zanasi, S., "Combinatorial Auction Based Virtual Machine Provisioning and Allocation In Cloud," (2013), *Wayne State University Dissertation Paper*, Paper 720.

14

© 2010 IBM Corporation

---

Introduction


[Second Price Auctions?](#)

---

Introduction


[SLAs](#)

---

Introduction


[SLAs](#)

---

Introduction


[SLAs](#)

- A provider optimizing capacity has a problem with the quality of the input

- Capacity is planned in practice solely based on the observed historical demand, error rate, and a forecast

- Forecast quality quickly deteriorates with:
  - Prediction horizon going further into the future
  - Historic data going back into the past

- Performance of forecasting often depends on fine tuning

15

© 2010 IBM Corporation

© 2010 IBM Corporation

Introduction



## Introduction

### Do we have SLAs today in an IaaS Cloud?



17

## Introduction

### IBM

### Do we have SLAs today in an IaaS Cloud?

**Finally, Real SLAs for Cloud Computing**

The SLA adopted for Cloud Servers™ is just as good as the one Rackspace provides for traditional hosts. It provides remedies for any downtime event caused by the network, data center, or infrastructure.

AWS will use commercially reasonable efforts to make EC2 available with an Annual Uptime Percentage of at least **99.95%** during the Service Year. If AWS fails to meet the Annual Uptime Percentage commitment, you will be eligible to receive a Service Credit as described below.

**rackspace® HOSTING**

**amazon web services**

**GO GRID**

**10,000% Guaranteed, 100% Uptime** Service Level Agreement

Customer's fees for the impacted Service Feature for the duration of the Failure. (For example, where applicable: a failure lasting seven (7) hours would result in credit of seven hundred (700) hours of free service for the feature in question)\*

**10,000% Service Credit\*** is a credit equivalent to **one hundred (100) times** supplements the Terms of Service and together such documents, and other **service level agreements** (SLAs) between the Customer and the Company.

**GO GRID**

## Introduction

### What do we **really** get? (assorted examples)

**rackspace® HOSTING**

We guarantee that our **data center network** will be available **100%** of the time in any given month (sliding period, excluding scheduled maintenance).

We guarantee that **data center HVAC and power** will be functioning **100%** of the time in any given month (sliding period, excluding scheduled maintenance, or heat problems).

We guarantee that function behavior in cloud environments will be **the same** as on-premises.

The minimum period of Failure eligible for a credit is 15 minutes, and shorter periods will not be aggregated. The maximum credit for any single Failure is **one month's Service fees**. In the event that multiple periods of Failure overlap in time, credits will not be aggregated, and Customer will receive credit only for the longest such period of Failure.

**amazon web services**

"Amazon guarantees the availability of the platform and the performance of the services it provides to the customer. Amazon guarantees that the platform and the services will be available 99.9% of the time in any given month (sliding period, excluding scheduled maintenance, or heat problems). We guarantee that function behavior in cloud environments will be the same as on-premises."

**GO GRID**

18

**Finally, Real SLAs for Cloud Computing**

The SLA adopted for Cloud Servers™ is just as good as the one Rackspace provides for traditional hosts. It provides remedies for any downtime event caused by the network, data center, or infrastructure.

AWS will use commercially reasonable efforts to make EC2 available with an Annual Uptime Percentage of at least **99.95%** during the Service Year. If AWS fails to meet the Annual Uptime Percentage commitment, you will be eligible to receive a Service Credit as described below.

**rackspace® HOSTING**

**amazon web services**

**GO GRID**

**10,000% Guaranteed, 100% Uptime** Service Level Agreement

Customer's fees for the impacted Service Feature for the duration of the Failure. (For example, where applicable: a failure lasting seven (7) hours would result in credit of seven hundred (700) hours of free service for the feature in question)\*

**10,000% Service Credit\*** is a credit equivalent to **one hundred (100) times** supplements the Terms of Service and together such documents, and other **service level agreements** (SLAs) between the Customer and the Company.

**GO GRID**

**Current Cloud SLA Practices Summary**

- Inflated promises
- Standard SLA with "one size fits all" availability SLO
- Obfuscation of provider commitments:
  - Being non-specific about maximum maintenance time per billing period
  - Being non-specific about maintenance head-up warning
  - Unavailability periods aggregation trickery
  - Requiring customers to understand failures on the vendor side
  - Being non-specific about availability tests to verify SLA compliance
  - Placing the onus of damage proof on the customer
- Refund policies that do not compensate the loss of value to the business due to unavailability

© 2010 IBM Corporation

© 2010 IBM Corporation

19

**Impact of downtime on business**

- Loss of profits
  - Impact on stock price
  - Loss of cash flow from debtors
  - Loss of customers (lifetime value of each)
  - Market share loss of product
  - Cost of fixing / replacing equipment
  - Cost of fixing / replacing software
  - Salaries paid to staff unable to undertake productive work
  - Salaries paid to staff to recover work backlog and maintain deadlines
  - Cost of re-creation and recovery of lost data
  - Interest value on deferred billings
  - Additional cost of credit through reduced credit rating
  - Fines and penalties for non-compliance
  - Liability claims
  - Additional cost of advertising, PR and marketing to reassure customers and prospects to retain market share
  - Additional cost of working ; administrative costs; travel, etc.
  - **Cloud:** *"we will give you a free service next month"*
- © 2010 IBM Corporation
- 21

**To remind us about our problem**

- Find minimal physical resources configuration to provide service to the customer, where the value of using the service worthy at least of what the customer pays for it
- SLAs are an approximation of the customer valuation (being arbitrarily far from the true consumer valuation)
- **Takeaway:** SLA construction is intimately related to capacity management
  - Short term (e.g., feedback loop control) resource management is done w.r.t. SLAs that might be suboptimal
- **In this talk: we don't discuss deriving optimal SLAs**
  - The focus will be on an easier problem: make SLAs **more specific** and optimize capacity w.r.t. them

© 2010 IBM Corporation© 2010 IBM Corporation21**...and there is also another input problem**

- Find minimal physical resources configuration to provide service to the customer, where the value of using the service worthy at least of what the customer pays for it
- Environment is very dynamic
  - There are long running services and short running services, tenants come and go, hardware changes, failures happen, disasters happen...
  - The environment changes as capacity optimization cycle completes...

21© 2010 IBM Corporation**...and, yeah, there is one more "small" input problem**

- Administrators don't like your tool monitoring their production environment:
  - Too much storage overhead
  - Too much network overhead
  - Too much disturbance to services normal operation

© 2010 IBM Corporation

So....

- No input on true customer valuations
- No stable VM populations
- More often than not no meaningful SLAs
- Scarce resources for capacity management

© 2010 IBM Corporation

25

Before moving forward with theory: **a disclaimer**

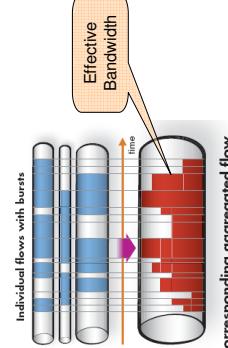
- **A distance between theory and practice is shorter in theory than in practice**
- Human administrator is the bottleneck for anything “smart”
- **In practice:** capacity management helps **slowing down** procurement cycle, but nobody [that I know] is looking for squeezing the last drop of optimality and/or accuracy

© 2010 IBM Corporation

26

Introduction  
Resource Demand

Introduction  
Cloud is cost-efficient thanks to Statistical Multiplexing



Corresponding aggregated flow

**Over-commit:** the total capacity of the shared resource is allowed to be *much* smaller than the maximum total demand

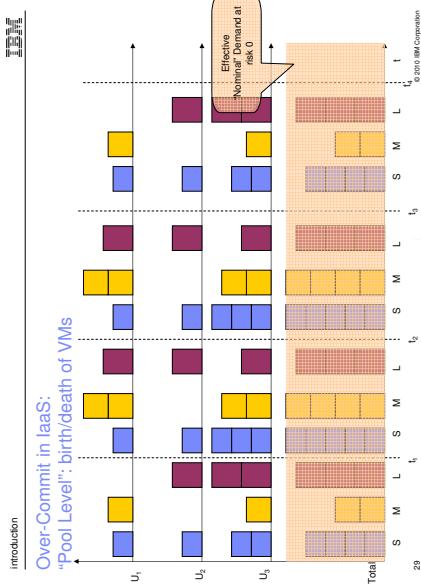
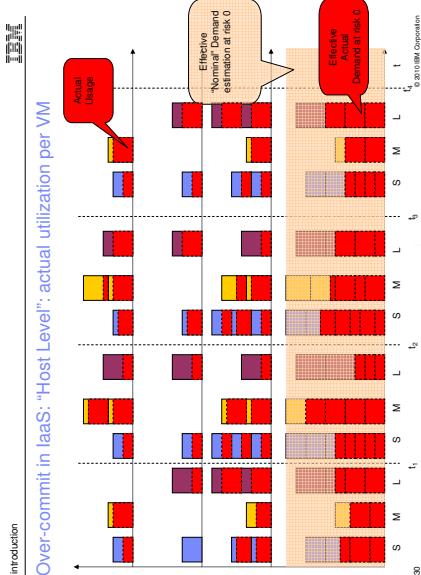
© 2010 IBM Corporation

28

- Can be expressed
  - In terms of “nominal allocations”
  - In terms of “actual utilization”

© 2010 IBM Corporation

29



- No statistics about actual usage of a VM are available
- High dynamics with VMs arriving and departing: no constant VM population
- A resource should be provided as an invisible "whole" (e.g., bare metal cloud server)

### Over-Commit on Actual Demand

- David Breitgand and Amir Epstein, "Improving Consolidation of Virtual Machines with Risk-aware Bandwidth Oversubscription in Compute Clouds", INFOCOM'12, March 25-30.

## IBM

---

### Standard SLA clause

results

- Standard availability clause in IaaS SLA:
- Percentile of the billing period when **pinging** VM succeeds



©2010 IBM Corporation

33

## IBM

---

### Extended Availability SLA

results

- Assume virtual hardware discrete types  $a_1, a_2, \dots, a_n$
- For simplicity, assume a single elasticity range for all services:  $R = [R_{\min}, R_{\max}]$
- Each service  $S_i$  is guaranteed to successfully assume configuration  $G_i = < a_1, a_2, \dots, a_n >$  subject to <elasticity range>, with probability  $\langle p \rangle$  computed over <billing period>
- $p \leq p_a$ , where  $p_a$  is the standard availability clause probability
- Rationale: guarantee on assuming desired configuration
- Interpretation: guarantee on probability to launch a VM of any type, subject to elasticity range

©2010 IBM Corporation

34

## IBM

---

### Critical observation

results

- We do not know distributions of  $Y_{ij}$
- Small contributions to  $X_i$
- For large clouds effect of dependencies diminishes
- Can be treated as independent
- Central Limit Theorem:  $X_i$  asymptotically converges to normal distribution
- Should also be identically distributed, but we ignore that in this work

©2010 IBM Corporation

35

## IBM

---

### SLA-aware cloud over-commit (CNSM'12)

results

#### Let $n$ be the total number of workloads (services)

#### Let $i$ be number of VM discrete types

- Let  $Y_{ij}$  be the random variable representing number of VM instances of type  $j$  used by workload  $j$
- $X_i = \sum_{j=1}^n Y_{ij}$  is the total number of VM instances of type  $i$  in the cloud

#### Definition 1: Nominal Demand is $\bar{X} = (X_1, X_2, \dots, X_l)$

#### Definition 2: Effective Nominal Demand

- Minimal vector  $D$ , such that

$$Pr(\bigwedge_i (X_i \leq D_i)) \geq p$$

- Union bound:  $Pr(V_i, X_i > D_i) \leq \sum_{i=1}^l Pr(X_i > D_i)$

p

©2010 IBM Corporation

36

## results

Not so fast ☺

- The variables are discrete
- Normal distribution is truncated normal

$$D_i = [\mu'_i + Z_{p_i} \sigma'_i]$$

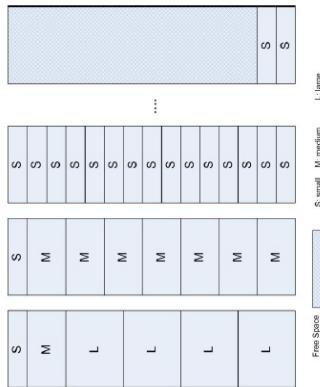
$$\mu'_i = \mu_i + \sigma_i \frac{\phi(-\frac{\mu_i}{\sigma_i})}{1 - \Phi(-\frac{\mu_i}{\sigma_i})} \quad \sigma'_i = \sigma_i \sqrt{1 - \frac{\phi(-\frac{\mu_i}{\sigma_i})}{1 - \Phi(-\frac{\mu_i}{\sigma_i})} \left( \frac{\phi(-\frac{\mu_i}{\sigma_i})}{1 - \Phi(-\frac{\mu_i}{\sigma_i})} + \frac{\mu_i}{\sigma_i} \right)}$$

37

© 2010 IBM Corporation

## results

Illustration: Sample VM slots reservation



38

© 2010 IBM Corporation

## Main Virtues

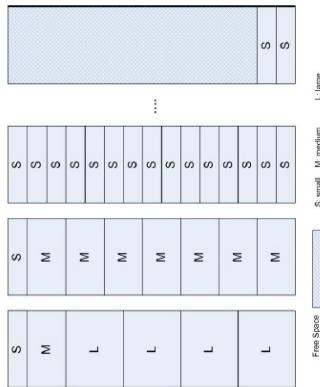
- Computationally and storage efficient
- Extremely simple
- Simple placement
- Explicit calculation of capacity based on risk perceived as tolerable → easy to transform into a policy
- Fast: easy to do “what-if?” analysis to extract providers’ true valuations of resource congestion

## A framework is more important than a specific algorithm

- If this is not the case, nominal strategy protects quality of experience (performance), but might be wasteful on resources

## results

Illustration: Sample VM slots reservation



39

© 2010 IBM Corporation

## results

Things missed by nominal strategy

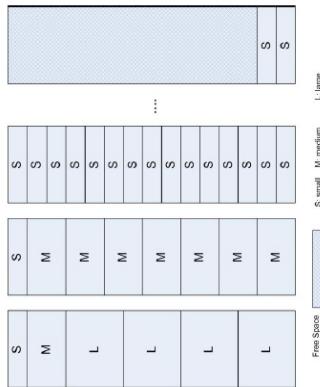
- Nominal capacity allocation strategy treats VMs/resources as indivisible “wholes”
- When should we expect it to produce best results?
- Gives best results when most of the time actual resource utilization of most of VMs on all of resource types is close to nominal discrete configuration of the VMs

39

© 2010 IBM Corporation

## results

Illustration: Sample VM slots reservation



40

© 2010 IBM Corporation



## Actual Utilization Over-Commit Strategy to the rescue!

- David Breitgand and Amir Epstein, "Improving Consolidation of Virtual Machines with Risk-aware Bandwidth Oversubscription in Compute Clouds", INFOCOM'12, March 25-30.

41

© 2010 IBM Corporation



## Stochastic Bin Packing Problem (SBP)

- $S = \{X_1, \dots, X_n\}$  – Set of items
  - $X_i$  – random variable representing the size (bandwidth demand) of item  $i$
  - $p$  – overflow probability
  - Goal: Partition the set  $S$  into the smallest number of subsets (bins)  $S_1, \dots, S_k$  such that
- $$\Pr[\sum_{i:X_i \in S_j} X_i > 1] \leq p \quad \text{for } 1 \leq j \leq k$$
- $p$  represents an SLA-stipulated value

42

© 2010 IBM Corporation



## Related Work – Bin Packing

- The problem is NP-hard
- Bin packing is hard to approximate to a factor better than  $\frac{3}{2}$  unless  $P=NP$ .
- First Fit Decreasing (FFD) has asymptotic approximation ratio of  $\frac{11}{9}$  and (absolute) approximation ratio of  $\frac{32}{27}$ .
- MFFD algorithm has asymptotic approximation ratio of  $\frac{71}{60}$ .
- AFPTAS exists.
- Online bin packing
  - First Fit (FF) has competitive ratio of  $\frac{17}{10}$ .
  - Best upper and lower bounds are  $1.58899$  and  $1.54014$ , respectively.

43

© 2010 IBM Corporation



## Related Work – Stochastic Bin Packing

- $O(\sqrt{\frac{\log p^{-1}}{\log \log p^{-1}}})$ -approximation for SBP with Bernoulli variables [Kleinberg et. al 1997]
- SBP with Poisson, Exponential and Bernoulli variables [Goei and Indik 1999]
  - PTAS exists for Poisson and exponential distributions.
  - Quasi-PTAS exists for Bernoulli variables.
  - These results relax bin capacity and overflow probability constraints by a factor  $1+\epsilon$ .
- $(1 + \sqrt{2})(1 + \epsilon)$  - competitive algorithm for SBP with normal variables [Wang et. al 2011]

44

© 2010 IBM Corporation

## Our Results (Breitgand and Epstein INFOCOM'12)

- 2-approximation algorithm for SBP with normal variables
- $(2+\epsilon)$ -competitive algorithm for online SBP with normal variables

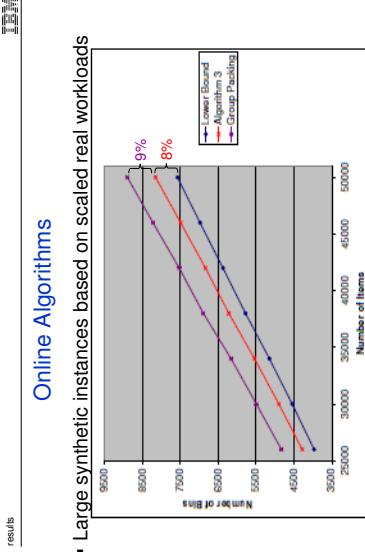
– Best known

## Intuition

- Collocating “bursty” items (VMs) together reduces effective size
- Normality assumption: relatively small number of VMs per host
  - For large hosts (e.g., large IBM Power machines), normality assumption can be dropped

© 2010 IBM Corporation

46



© 2010 IBM Corporation

47

$p$	Algorithm 3 (Online)	Algorithm 1 (Approx.)	Group Packing	FID	FF	Algorithm 2 (LB)
0.1	164	46	505	332	334	144
0.01	215	95	785	519	522	192
0.001	263	243	881	656	662	237

© 2010 IBM Corporation

48

IBM

### Summary of the methodology

- Make SLA meaningful – starting point
- Calculate effective capacity with respect to observed demand and target resource congestion probability
- Calculate placement for effective capacity
- Continuously update effective capacity
- Recalculate placement only when a significant (affecting target congestion probability) is spotted

49

© 2010 IBM Corporation

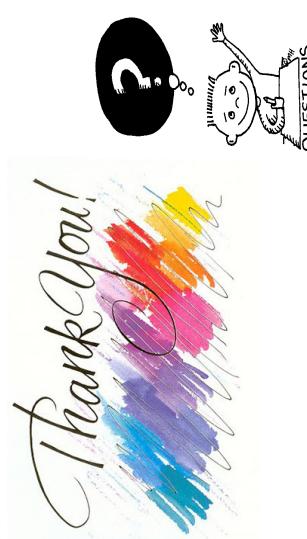
IBM

### Conclusions & Future Research: where does this become really hard?

- Placement Constraints: cannot use simple BP anymore
  - David Breslau and Amit Epstein, "SLA-aware placement of multi-virtual machine elastic services in compute clouds", IFIP/IEEE Integrated Network Management (IM'11), pp. 161-168, Dublin, Ireland
  - Extends: B. Ugoniak, A. L. Rosenberg, and J. J. Stanoić, "Application placement on a cluster of servers," *Int. J. Found. Comput. Sci.*, vol. 18, no. 5, pp. 1023-1041, 2007.
  - Elastic Services Placement Problem (ESPP) generalizes GAP, but does not admit constant factor approximation  $n^{1/2-\epsilon}$  for any constant  $\epsilon > 0$
- Performance degradation due to non-virtualized resources
  - L2 cache? Bus? Not visible metrics
- **Personal appreciation 1:** these two problems are the core ones impeding progress on **systematically** improving cost-efficiency in IaaS
- **Personal appreciation 2:** solution is likely in apps so collaborating with infrastructure provider

© 2010 IBM Corporation

IBM



51

© 2010 IBM Corporation

## DEADLINE SCHEDULING FOR BIG-DATA JOBS AND FAULT-TOLERANCE OF DATACENTER APPLICATIONS

Peter Bodik, Microsoft Research

I will describe two projects in the space of resource allocation in large-scale datacenters: Jockey -- scheduling big-data jobs to meet latency deadlines and application placement in datacenters to survive large-scale hardware failures.

Many big-data jobs, running in Hadoop MapReduce or Microsoft's Cosmos, require completion by a certain deadline. Missing a deadline might lead to reduced productivity of data analysts, stale content presented by a search engine, or even a financial penalty. However, today's cluster schedulers do not support specifying a deadline for a job and provide no guarantees on job completion. I will describe Jockey, a framework for providing deadline guarantees for big-data jobs. Offline, Jockey uses past executions of a job to build a model of the job and then, during job execution, uses the model in a control loop to adjust job resources to meet the specified deadline.

In the second half of the talk, I will talk about improving fault tolerance of applications deployed in datacenters. Datacenter networks have been designed to tolerate failures of network equipment and provide sufficient bandwidth. In practice, however, failures and maintenance of networking and power equipment often make tens to thousands of servers unavailable, and network congestion can increase service latency. Unfortunately, there exists an inherent tradeoff between achieving high fault tolerance for applications deployed in a datacenter and reducing bandwidth usage in network core. Spreading servers across fault domains improves fault tolerance, but requires additional bandwidth, while deploying servers together reduces bandwidth usage, but also decreases fault tolerance. We present a detailed analysis of a large-scale Web application and its communication patterns. Based on that, we propose and evaluate a novel optimization framework that achieves both high fault tolerance and significantly reduces bandwidth usage in the network core by exploiting the skewness in the observed communication patterns.



## Why care about deadlines for big-data jobs?

### Big-data job deadlines

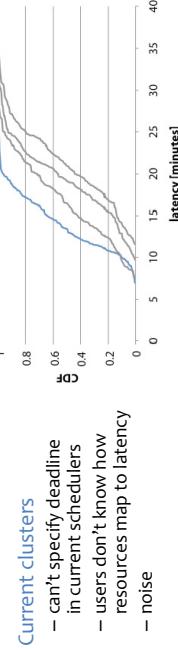
+

### Fault-tolerant resource allocation

Peter Bodik

Microsoft Research

- Important big-data jobs have to finish on time
- missed deadline = delayed updates on site, financial penalty, productivity



## Jockey: meeting deadlines for big-data jobs

### Cosmos 101

- big-data platform in Microsoft
- job = SQL query + user C# code
- job compiled/optimized using SQL-like optimizer to a DAG of stages/vertices
- big jobs have 100s of stages, 1M vertices

### Jockey

- input: single job with a deadline, past job runs
- offline: builds a job model
- run time: control loop adjusts allocation

## Job model = past job runs + simulation

### Job model

- input: current progress, allocation
- output: remaining time to completion

### Example

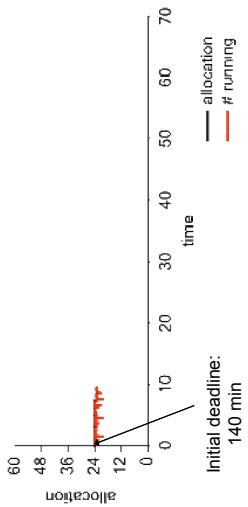
- deadline = 30 min
- after 10 min, completed 50%
- will set allocation to 30 tokens

### Issues

- in practice need to trade off between many jobs

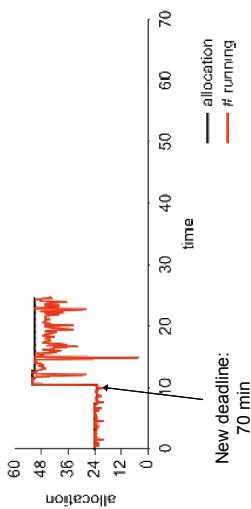
	10 tokens	20 tokens	30 tokens
10%	60 min	40 min	25 min
20%	59 min	39 min	24 min
30%	58 min	37 min	22 min
40%	56 min	36 min	21 min
50%	54 min	34 min	20 min

## Jockey in Action



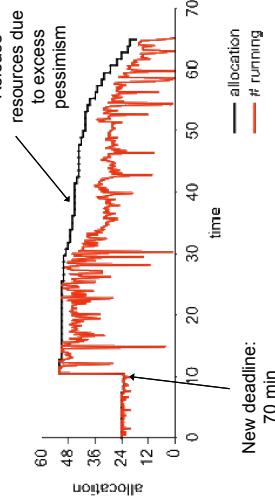
5

## Jockey in Action



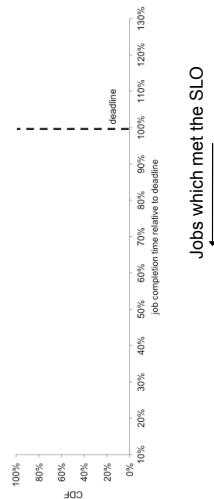
6

## Jockey in Action

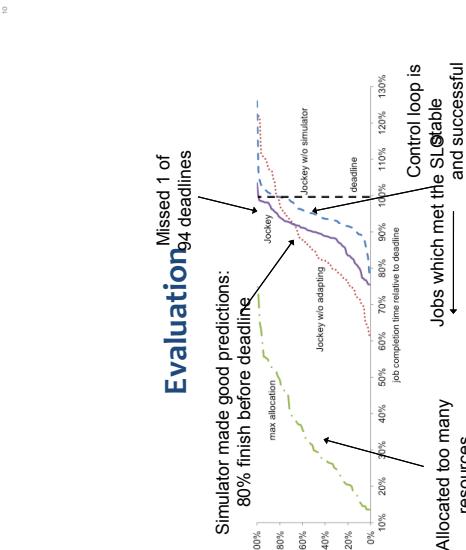
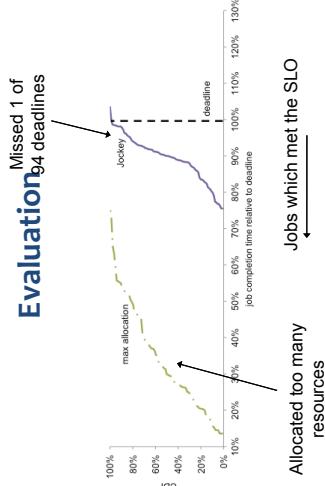
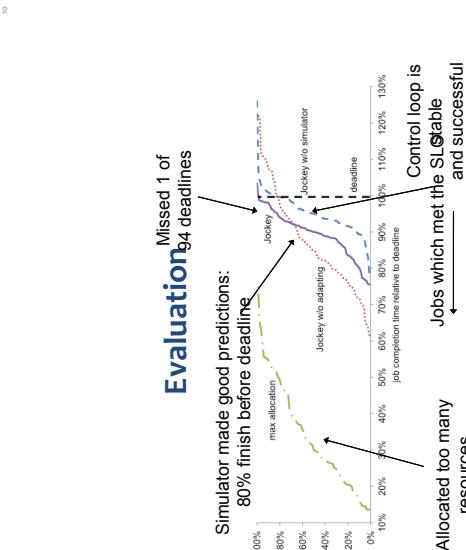
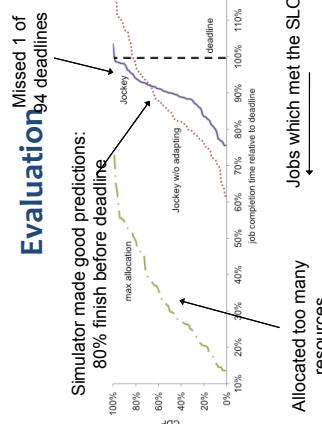
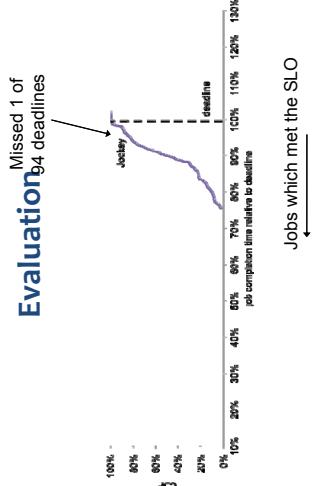


7

## Evaluation



8



## What's missing?

- multiple jobs/deadlines, multiple pipelines
- better representation of job state

## FAULT-TOLERANT RESOURCE ALLOCATION

### How to allocate services to physical machines?

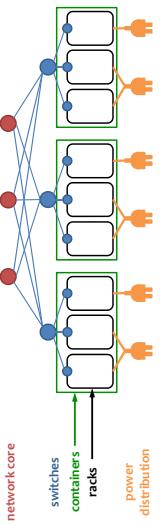


Three important metrics considered together

- FT: service fault tolerance
- BW: bandwidth usage
- #N: # machine moves to reach target allocation

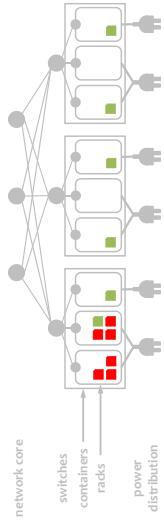
SIGCOMM 2012, Surviving Failures in Bandwidth-Constrained Datacenters  
Peter Bodik, Shai Menache, Mosharaf Chowdhury, Pradeepkumar Mani, David A. Maltz, and Ion Stoica

### FT: Improving fault tolerance of software services



- Complex fault domains: networking, power, cooling
- Worst-case survival = fraction of service available during single worst-case failure
  - corresponds to service throughput during failure

## FT: Service allocation impacts worst-case survival

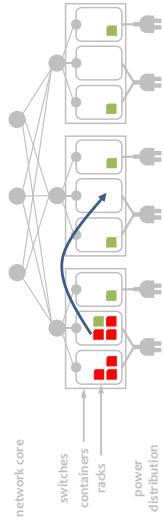


**Worst-case survival:**

- red service: 0% – same container, power
- green service: 67% – different containers, power

17

## #M: Need incremental allocation algorithms

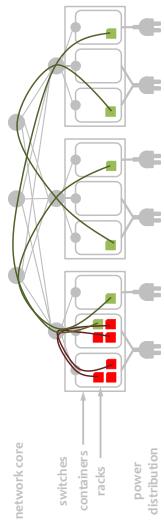


**High cost of machine move**

- need to deploy potentially TB of data
- warm up caches
- could take tens of min, impact network

19

## BW: Reduce bandwidth usage on constrained links



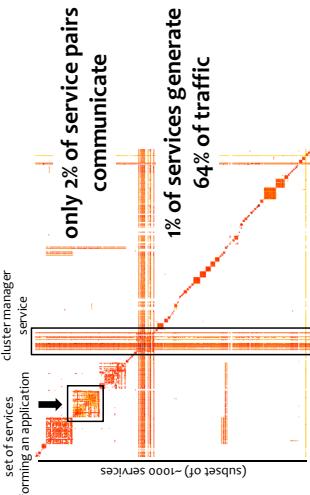
**BW = bandwidth usage in the core**

**Goal**

- reduce cost of infrastructure
- consider other service location constraints

18

## Service communication matrix is very sparse and skewed



20

## Formulate as convex optimization

Spread machines across all fault domains

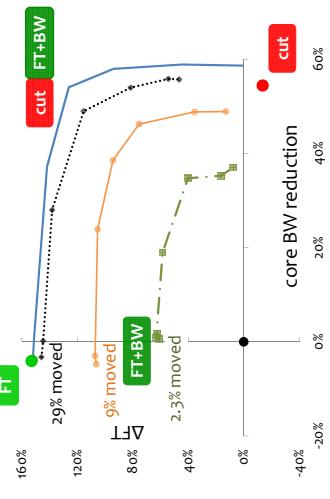
$$\min \alpha BW + \sum_s c_s \sum_f w_f \cdot z_{s,f}^2$$

number of machines  
of services in domain  $f$   
fault domain weight  
service weight

Advantages of convex cost function

- local actions (machine swaps) lead to improvement of global metric
- directly considers #M

## Evaluation



22

## Potential future work

### Deadline scheduling

- multiple deadline jobs (with different penalties for missing deadline)
- dependencies across jobs
- maximize total utility
- adapt to new jobs arriving, reduces capacity, ...

### Resource allocation

- consider structure of the applications, e.g. data partitions and replications
- dependencies across services

23

**CONTROL ISSUES IN WAREHOUSE-SCALE DATACENTERS****John Wilkes, Google**

Google's compute and storage clusters are managed by a raft of different software systems that attempt to achieve multiple, partially conflicting, goals simultaneously. I will present an overview of some of these systems, and highlight some of the challenges that we're facing along the way, on the grounds that many challenges also represent good research opportunities.



# Control issues in warehouse-scale datacenters



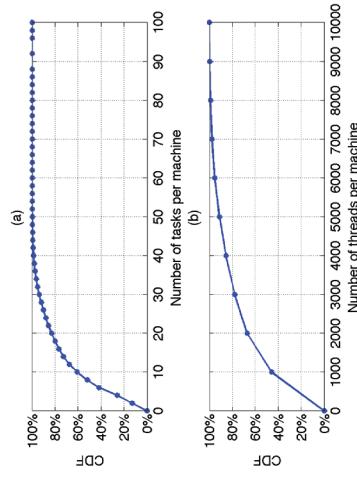
john wilkes  
Cloud Control Workshop, Lund, Sweden  
May 2014



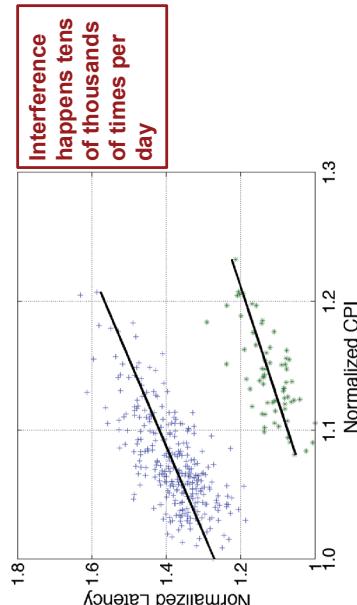
**The problem**  
high utilization => resource sharing

**The problem**  
resource sharing => interference

From: CPM: CPU performance isolation for shared compute clusters, EuroSys'13.



Interference  
happens tens  
of thousands  
of times per  
day

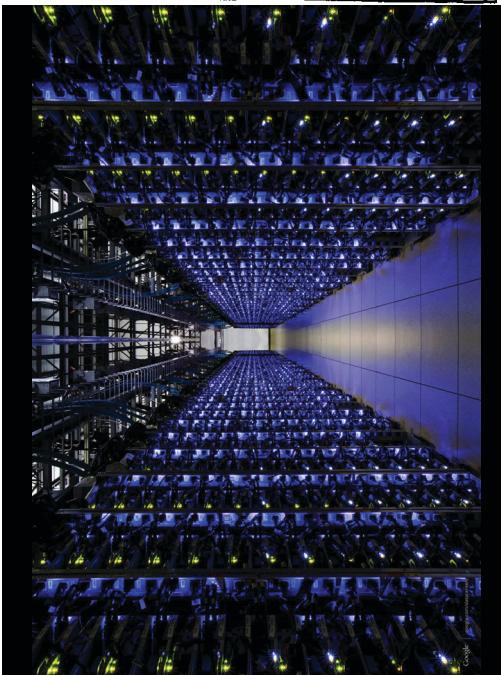
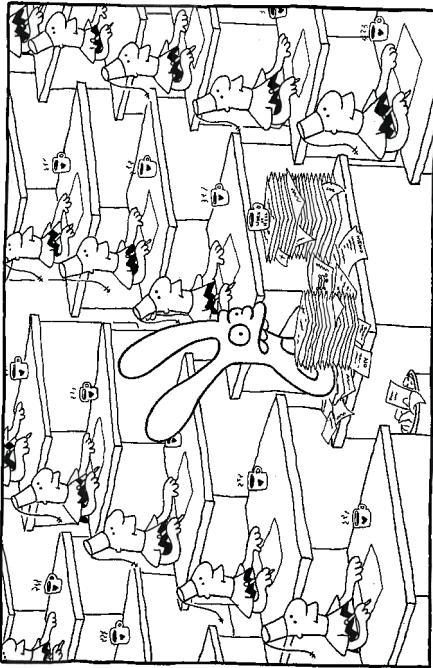
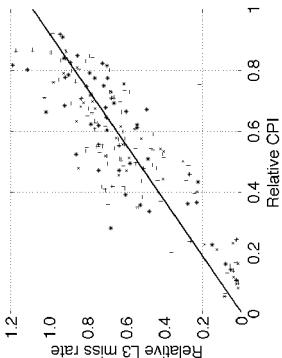


## Our solution: CPI<sup>2</sup>

a simple control system

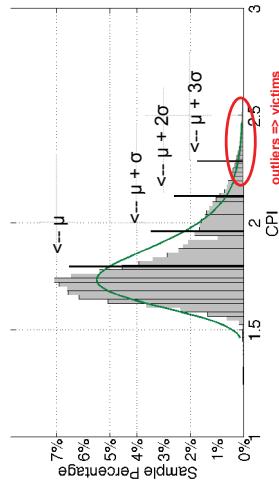
## Why use CPI?

1. Monitor Cycles Per Instruction (CPI)
  2. Learn anomalous behaviors
  3. Identify a likely antagonist
  4. Throttle it to shield victims
- It's cheap: < 0.1% CPU overhead, invisible to users
  - It's stable (across time and space)
  - It correlates well with L3 cache miss rate

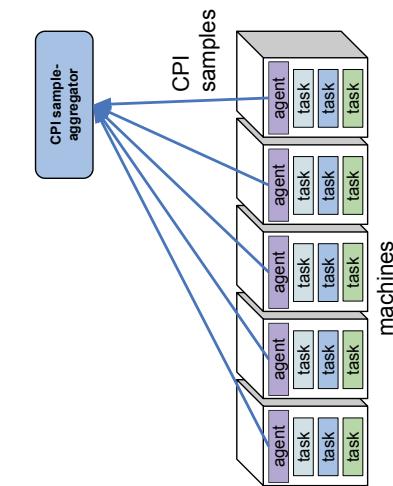


## Gathering CPI

- Build a CPI profile for a job
- per-cluster, per-platform
- mean ( $\mu$ ) & stddev ( $\sigma$ )

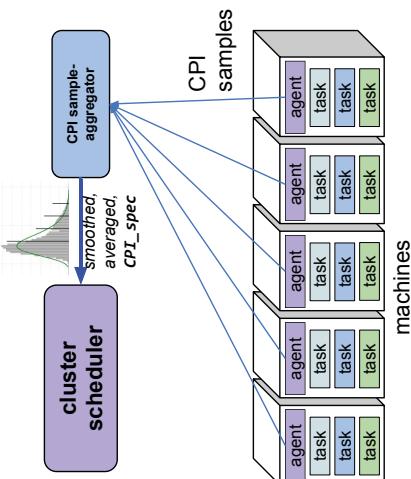
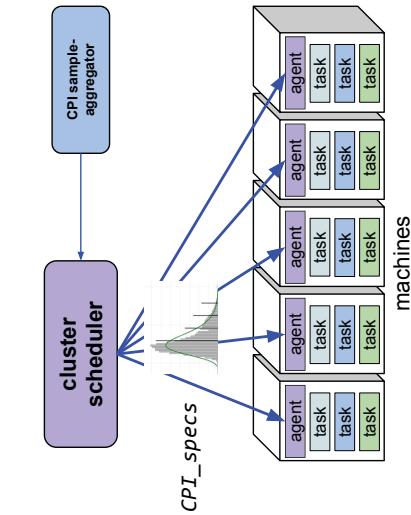


## Gathering CPI



## Gathering CPI

### Using CPI to detect an anomaly



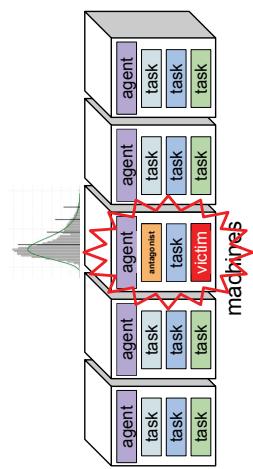
## Using CPI to detect an anomaly

### Now what?

Goal: reduce the effect of the antagonist

Let's **throttle** the antagonist!

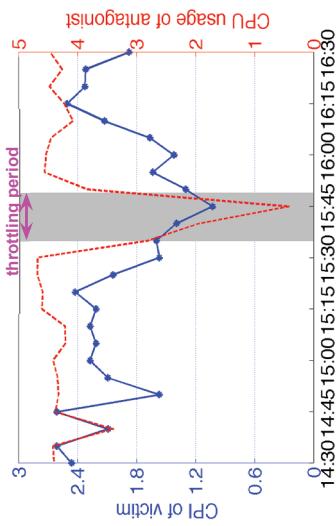
- CPU hard-capping: 0.1 core for 5 minutes



Restrictions:

- only throttle batch jobs
- only help “important” victims

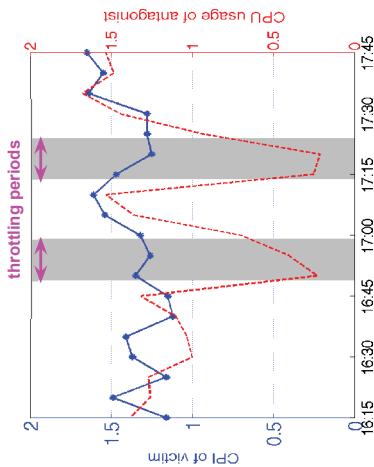
### A motivating example



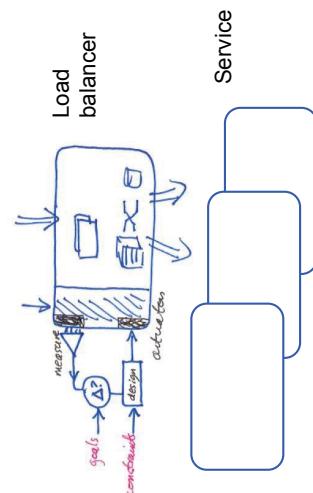
**What could possibly go wrong?**

## A not so good example

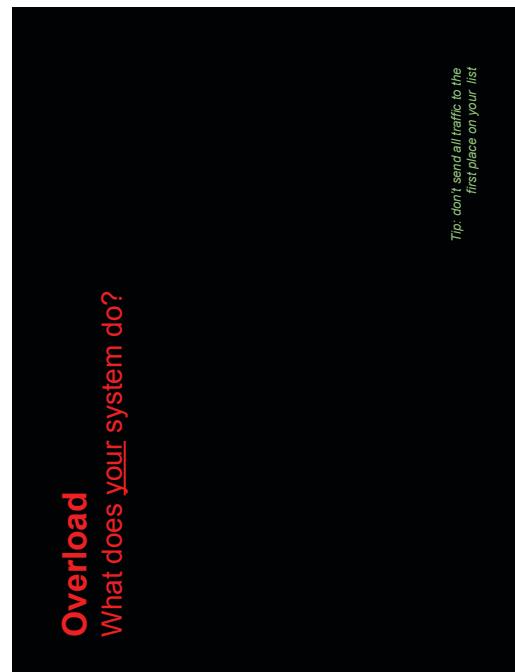
**Maybe batch-only was a bad idea?**  
After all: LS tasks have *load balancing*



**Maybe batch-only was a bad idea?**  
After all: LS tasks have *load balancing*



**Overload**  
What does your system do?



## Maybe batch-only was a bad idea?

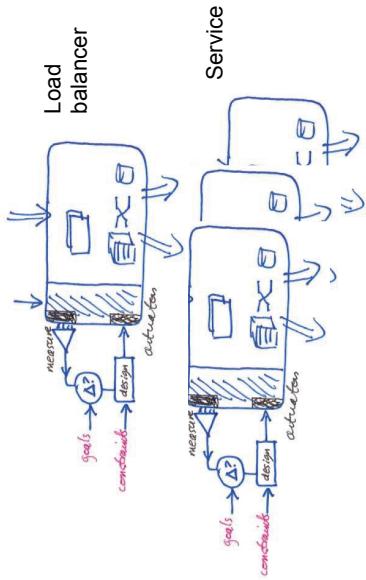
After all: LS tasks have *load balancing*

### Cascading failures

1. Overload-induced outage
  - o busy cluster => oops
2. No worries! Shunt load elsewhere!
  - o busy cluster => much oops (repeat)
  - o e.g., Gmail outage, 2009-02-24

## Maybe batch-only was a bad idea?

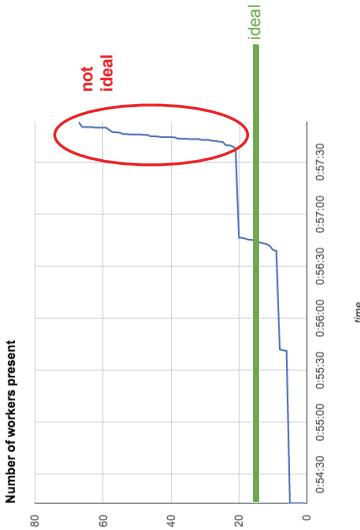
After all: LS tasks have *load balancing*



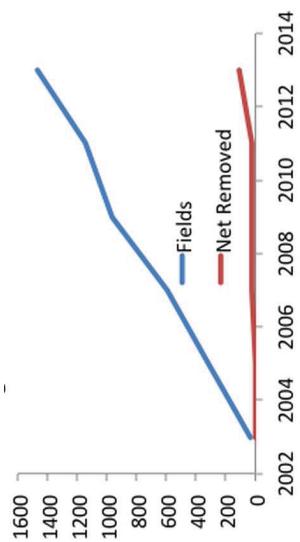
## Interacting control loops

1. Load-placement
  - few-second response times
2. Number-of-workers
  - few tens-of-seconds response times
3. Add a little signalling delay ...

## Auto-scaling to meet a job deadline



**No worries!**  
Just add a few more knobs ...



## Upload malformed configuration What does your system do?

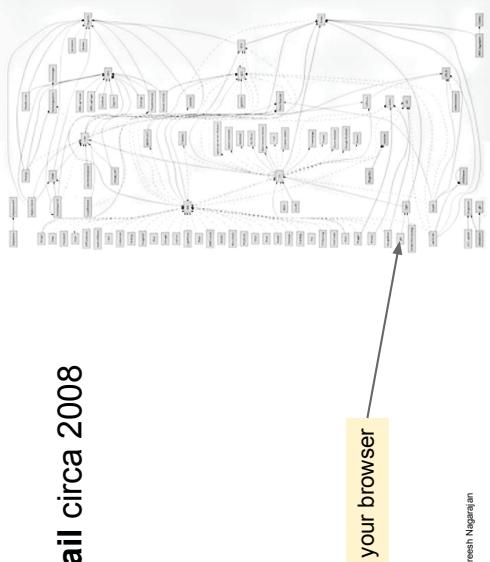
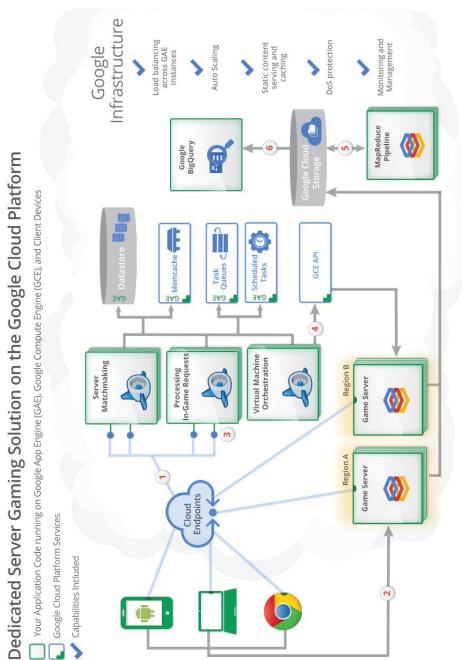
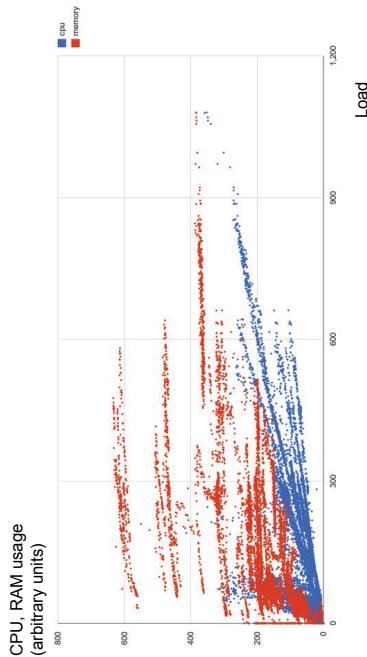


Image source: Hareesh Nagarajan



## Model building is hard



Maybe more monitoring would help?

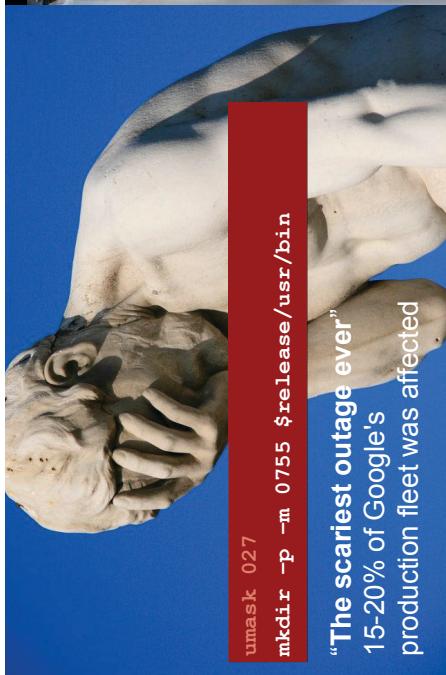
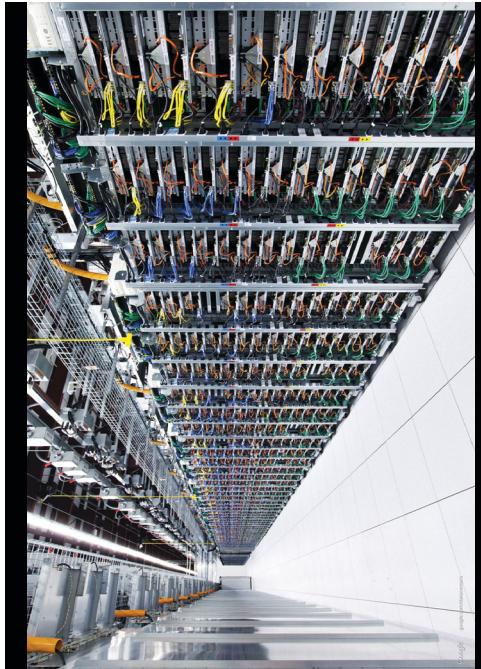
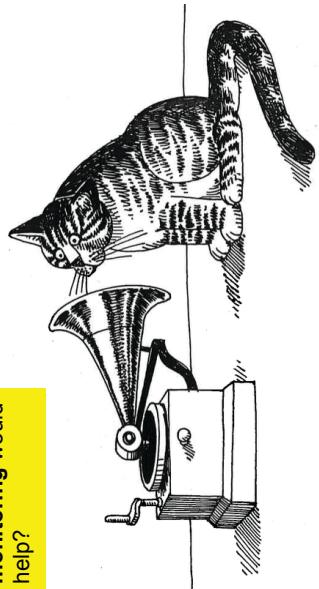


Photo credit: Alex E. Proimos | Creative Commons

## Delegation is hard

be careful what you ask for



**It's 3am and your pager goes off**

- are we in trouble?
  - are we about to get into trouble?
- **what should you do about it?**

## Summary

Control systems do not run in isolation

1. Do no harm
2. Make things better
3. **Assume the world is out to get you**  
“any sufficiently advanced incompetence is indistinguishable from malice”  
-- Grey's Law

**APPLICATION PERFORMANCE MANAGEMENT IN THE CLOUD USING LEARNING, OPTIMIZATION, AND CONTROL**  
**Xiaoyun Zhu, VMware Inc.**

Many businesses and organizations are increasingly relying on cloud based infrastructures and platforms to deliver their business-critical applications. In the meantime, a recent study shows that 79% of companies are concerned about the hidden costs of cloud services for their applications, citing “poor end-user experience due to performance bottlenecks” as their top management concern in relationship to cloud services. Existing practices in application performance management rely heavily on white-box modeling or heuristics-based, manual diagnostic approaches to find potential bottlenecks and remediation steps. However, the scalability and adaptivity of such approaches remain severely constrained, especially in a highly-dynamic, consolidated cloud environment. These challenges present unique opportunities in applying statistical learning, control, and optimization based techniques to developing model-based, automated application performance management frameworks. There has been a large body of research in this area in the last several years, but many problems remain. In this talk, I’ll highlight some of the performance and resource management techniques we have developed within VMware, along with related technical challenges, and discuss open research problems, in hope to attract more innovative ideas and solutions from a larger community of researchers and developers.

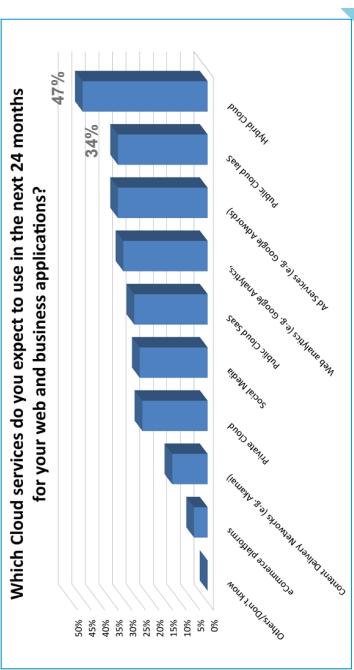


# Application Performance Management in the Cloud using Learning, Optimization, and Control

Xiaoyun Zhu

May 9, 2014

## Rising adoption of cloud-based services

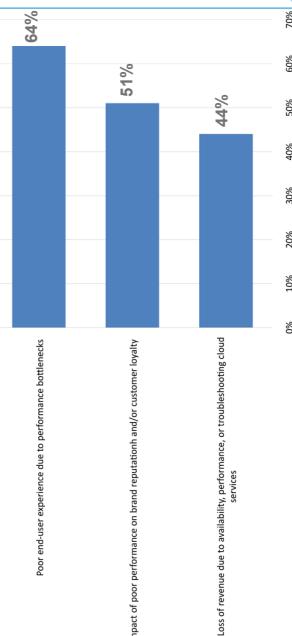


vmware

© 2014 VMware Inc. All rights reserved.

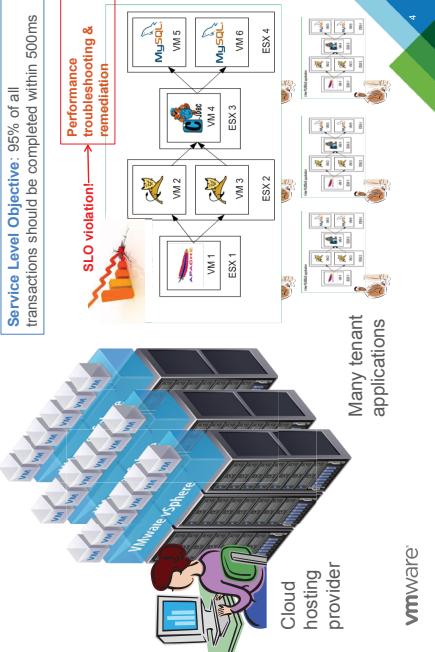
## Application performance – a real concern

What are your biggest concerns about managing  
Cloud services?



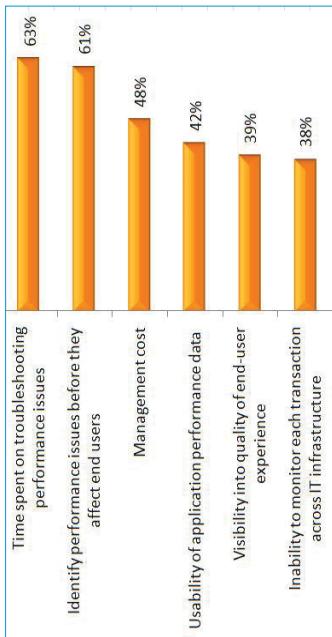
vmware

## Application performance management is hard



vmware

## Challenges in managing application performance



\* On average, **46.2 hours** spend in "war-room" scenarios each month

Source: Improving the usability of APM tools: essential capabilities and benefits. TRAC Research, June 2012, based on survey data from 400 IT organizations worldwide

**vmware**

## Infrastructure-level performance monitoring

### Physical host metrics

- System-level stats collected by the hypervisor
  - e.g., `esxtop` – CPU, memory, disk, network, interrupt
- CPU stats
  - %USED, %RUN, %RDY, %SYS, %OVRLP, %CSTRP, %WAIT, %IDLE, %SWPWT
- ~100s-1000s metrics per host!

### VM metrics

- Resource usage stats collected by the guest OS
  - e.g., `dsstat`, `iostat`
- ~10s metrics per VM
- Widely available on most platforms
- Available at a time scale of seconds to minutes

**vmware**

## APM-related problems we're working on

- Real-time performance monitoring
  - Infrastructure-level vs. application-level monitoring
- Automated performance modeling
  - Knowledge-driven vs. data-driven
  - Linear vs. nonlinear models
  - Offline vs. online modeling
- Computer-assisted performance troubleshooting
  - Correlation & model based problem localization
- Service level remediation via auto-scaling
  - Horizontal vs. vertical scaling



## Application-level performance monitoring

### Metrics reflecting end user experience

- Response times
- Throughput (or errors such as timed out requests)

### VMware Hyperic monitoring tool

- Agents deployed in VMs
- Auto-discovers types of applications running
- Plugins to extract application-related performance stats
  - Stats available at a time scale of minutes
  - Stats aggregated in Hyperic server
  - Supports over 80 different application components
  - Extensible framework to allow customized plugins

**vmware**



## APM-generated big data

- “APM tools were part of the huge **explosion in metric collection**, generating thousands of KPIs per application.”
- “83% of respondents agree that metric data collection has **grown >300%** in the last 4 years alone.”
- “88% of companies are only able to analyze **less than half** of the metric data they collect... **45%** analyze **less than a quarter** of the data.”
- “77% of respondents cannot effectively **correlate** business, customer experience, and IT metrics.”

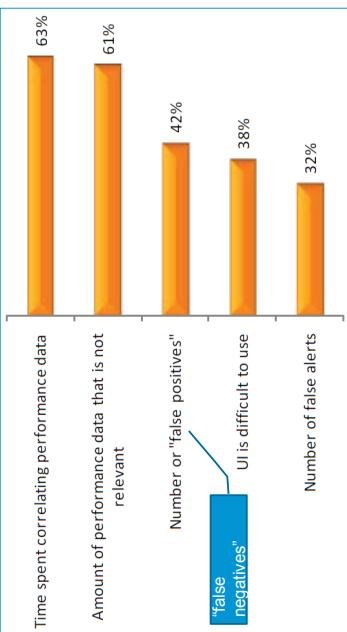
Source: “APM-generated big data boom.” Netwrix & APMDigest, July 2012, based on survey of US & UK IT professionals.

vmware

Source: Improving the usability of APM data: Essential capabilities and benefits. TRAC Research, June 2012, based on survey data from 400 IT organizations worldwide.

vmware

## Challenges in usability of performance data

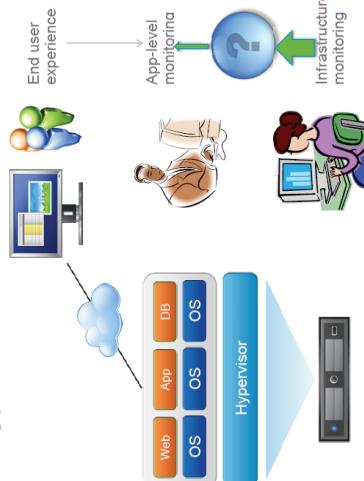


Source: Improving the usability of APM data: Essential capabilities and benefits. TRAC Research, June 2012, based on survey data from 400 IT organizations worldwide.

vmware

## The Semantic Gap challenge

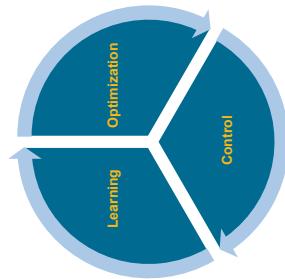
Correlating performance data from different sources



vmware

## Better IT analytics for APM automation

Three-pronged approach



vmware

12

11

## Semantic gap filled by performance models Learning-based approach

Traditional models harder to apply

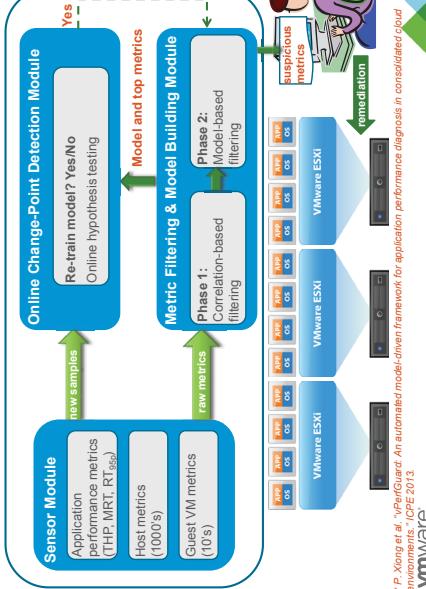
- First-principle models: Only exist for special cases (e.g., flow models)
- Queuing models: More suitable for aggregate/average behavior
- Architectural models: Require domain knowledge, harder to automate

Empirical models via statistical learning

- Data driven, easier to automate and scale
- Offline modeling** usually insufficient
  - Time-varying workloads
  - Changing system/software configurations
- Online modeling** (models updated on demand)
  - Need to be low overhead and adaptive

VMware

## Correlation and model based metric selection



\*P. Xiong et al., "PerGuard: An automated model-driven framework for application performance diagnosis in consolidated cloud environments," CPE 2013.

## Three key questions

• Q1: Which variables go into the model?

- Which system resources or parameters affect application performance the most?
- Correlation-based analysis to provide hints

• Q2: What kind of model should we use?

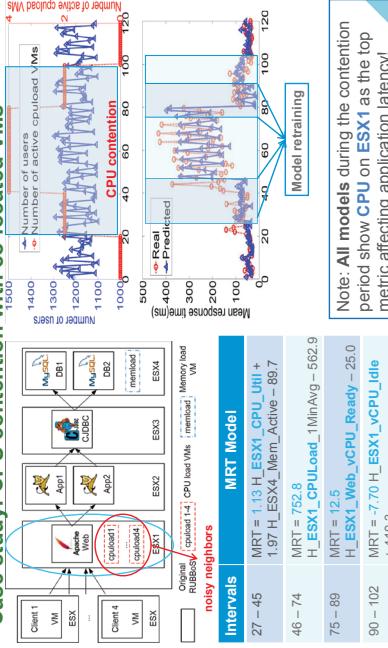
- Nonlinear models - better accuracy in general
- Linear regression models - cheaper to compute and easier to interpret

• Q3: How do we know our model is (still) accurate?

- Online change-point detection

VMware

## Case study: CPU contention with co-located VMs



VMware

Note: All models during the contention period show **CPU** on **ESX1** as the top metric affecting application latency!

VMware

16

14

13

12

11

10

9

8

7

6

5

4

3

2

1

0

15

16

## Challenges to ensure application performance

- Enterprise applications are **distributed** or **multi-tiered resources**
  - HW: CPU, memory, cache, network, storage
  - SW: threads, connection pool, locks
- **Time-varying** application behavior
  - Dynamic and **bursty** workload demands
  - Performance **interference** among co-hosted applications

vmware

vmware

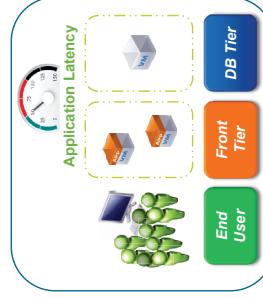
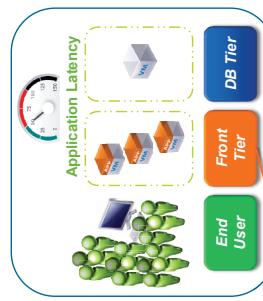
vmware

vmware

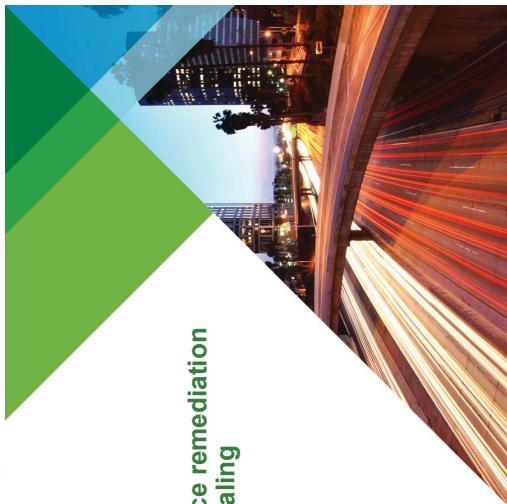
vmware

vmware

## Auto-Scaling to Maintain Application SLO



## Performance remediation via auto-scaling



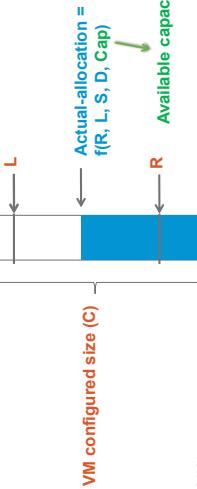
## Auto-Scaling to Maintain Application SLO



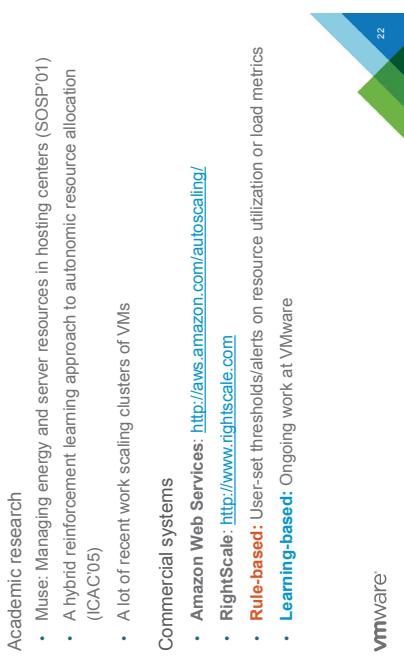
## Vertical scaling of resource containers

### Method 1: Dynamic resource control settings

- Available on various virtualization platforms
- For shared CPU, memory, disk I/O\*, network I/O\*:
  - Reservation ( $R$ )** – minimum guaranteed amount of resources
  - Limit ( $L$ )** – upper bound on resource consumption (non-work-conserving)
  - Shares ( $S$ )** – relative priority during resource contention
- VM's CPU/memory **demand ( $D$ )**: estimated by hypervisor, critical to actual allocation



## Horizontal scaling of applications



## Vertical scaling of resource containers

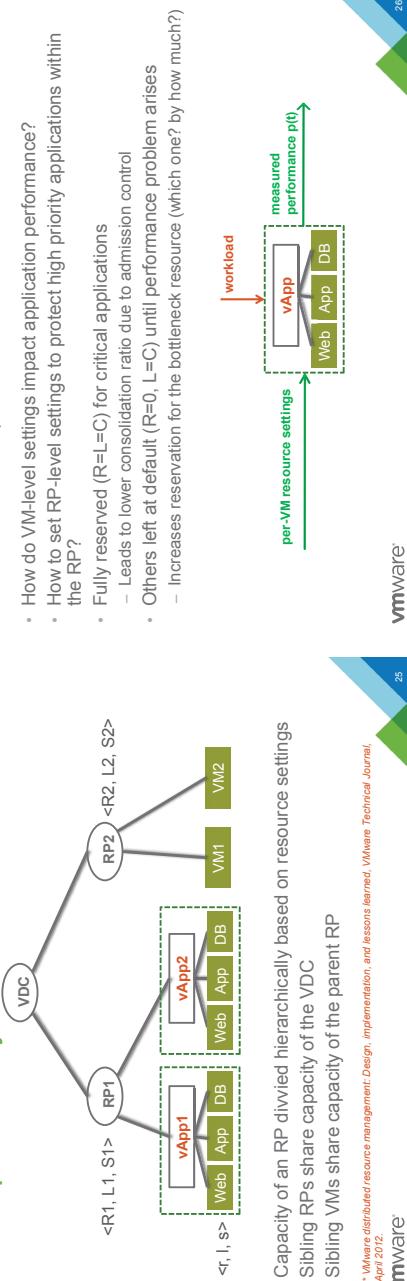
### Related work (not exhaustive)

- Tuning resource **limits** (aka. **caps**)
  - Adaptive control of virtualized resources in utility computing environments (Eurosys'07)
  - Autonomic resource management in virtualized data centers using fuzzy-logic-based applications (Cluster Computing Journal 2008)
  - Memory overbooking and dynamic control for Xen virtual machines in consolidated environment (IM09, memory limit)
  - Vertical scaling of prioritized VMs provisioning (GC'12)
  - Agile: Elastic distributed resources scaling for infrastructure-as-a-service (ICAC'13)
- Tuning resource **shares** (aka. **weights**)
  - Maximizing server utilization while meeting critical SLAs via weight-based collocation management (IM'13)
  - Application-driven dynamic vertical scaling of virtual machines in resource pools (NOMS'14)



## DRS (Distributed Resource Scheduler)

### Resource pool hierarchy



25

vmware

## Powerful knobs, hard to use

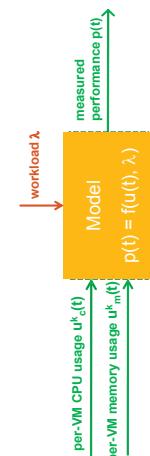
- How do VM-level settings impact application performance?
  - How to set RP-level settings to protect high priority applications within the RP?
    - Fully Reserved ( $R=L=C$ ) for critical applications
      - Leads to lower consolidation ratio due to admission control
      - Others left at default ( $R=0, L=C$ ) until performance problem arises
      - Increases reservation for the bottleneck resource (which one? by how much?)

26

vmware

## Performance model learned for each vApp

- Maps VM-level resource allocations to app-level performance
- Captures multiple tiers and multiple resource types
  - Choose a linear regression model (easy to compute)
  - Workload indirectly captured in model parameters
  - Model parameters updated online in each interval (tracks nonlinearity)



27

vmware

## Rule-based vs. model-based feedback control

Rule-based	Model-based
often involves no analytical model	requires an analytical model
driven by intuition and domain knowledge	driven by quantitative relationships
hard to control multiple knobs at the same time	captures interactions between multiple metrics
no concern of dynamics	considers dynamics and transient responses
threshold and heuristics based	standard control methods as building blocks
no systematic consideration of stability	systematically handles tradeoff between stability & performance

28

vmware

## Use optimization to handle design tradeoff

- An example cost function

$$J(\mathbf{u}(t+1)) = p(t+1) - p_{SLO} + \beta \|\mathbf{u}(t+1) - \mathbf{u}(t)\|^2$$

control cost

Tradeoff between performance and stability

Performance cost

control cost

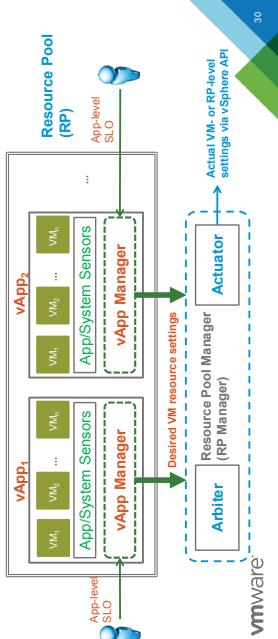
- Solve for optimal resource allocations

$$\mathbf{u}^*(t+1) = g(p(t), p_{SLO}, \mathbf{u}(t), \lambda, \beta)$$

vmware

## AppRM: Model-based vertical scaling

- Auto-tunes VM-level and RP-level resource control settings to meet application SLOs
  - For each application, **vApp Manager** translates its SLO into **desired** resource control settings at individual VM level
  - For each resource pool, **RP Manager** computes the **actual** VM- and RP-level resource settings to satisfy all critical applications

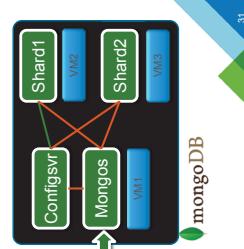


vmware

29

## Performance evaluation

- Application
  - MongoDB** – distributed data processing application with sharding
  - Rain – workload generation tool to generate dynamic workload
- Workload
  - Number of clients
  - Read/write mix
- Evaluation questions
  - Can the vApp Manager meet individual application SLO?
  - Can the RP Manager meet SLOs of multiple applications?

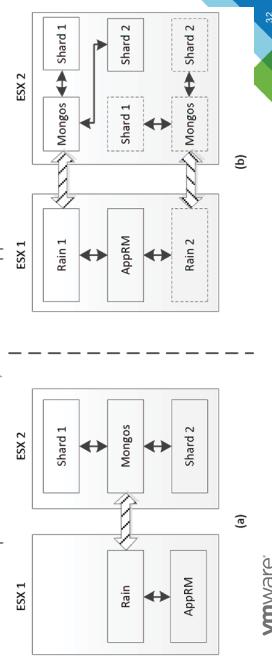


vmware

31

## Testbed setup

- Two ESX 5.0 GA hosts
  - ESX2 (12 cores, 96 GB) to emulate the capacity of a VDC
  - Three VMs per MongoDB instance (2 vCPUs, 4 GB)
  - One VM per instance of Rain, one VM for AppRM

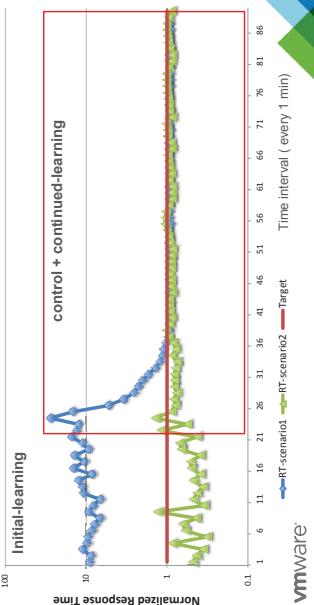


vmware

32

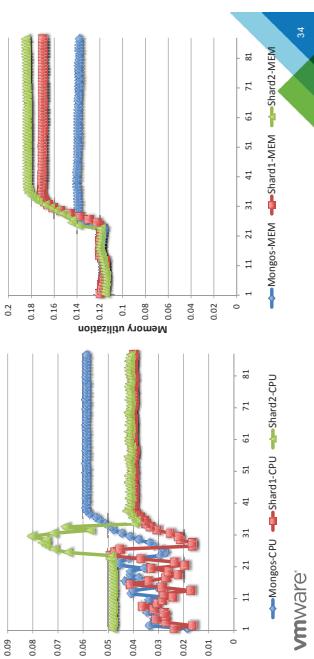
## Result: Meeting mean response time target

- Scenario1 - Initial settings: R = 0, L = 512 (MHz, MB)
- Scenario2 - Initial settings: R = 0, L = unlimited (cpu, mem)



## Resource utilization (under-provisioned case)

- Target response time = 300 ms
- Initial setting R = 0, L = 512 MHz/MB (under-provisioned)
- Initial setting R = 0, L = unlimited (cpu, mem)



## Vertical scaling of resource containers

### Method 2: Runtime reconfiguration of VM sizes

- Configured size for a VM

- #vCPUs
- Memory size
- #virtual disks, disk sizes
- #NICs

- ESX allows over-commitment of CPU and memory

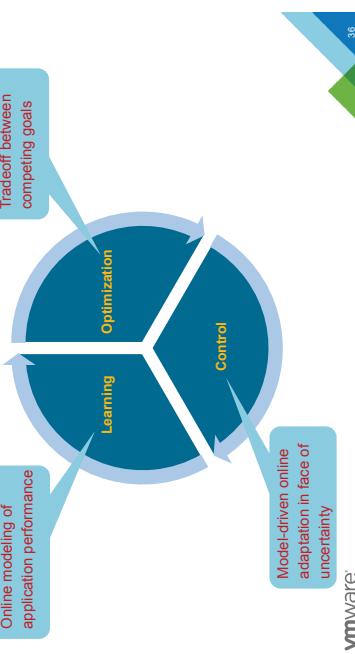
- $\text{Sum}(\text{VM-size}) \geq \text{host-capacity}$
- CPU/memory **Hot-add** supported by most recent OS's
  - Can be used to scale up a VM at runtime – work-in-progress
  - Need application support to leverage additional resources
- CPU/memory **Hot-remove** unsupported by most OS's
  - Requires VM reboot (undesirable)

VMware

## Recap: APM automation requires better analytics

- Online modeling of application performance

- Tradeoff between competing goals



VMware

## References

- X. Zhu, et al. "What does control theory bring to systems research?" *ACM SIGOPS Operating Systems Review*, 43(1), January 2009.
- P. Padala et al. "Automated control of multiple virtualized resources." *Eurosys 2009*.
- A. Gulati et al. "Cloud scale resource management: Challenges and techniques." *HotCloud 2011*.
- A. Gulati et al. "VMware distributed resource management: Design, implementation, and lessons learned." *VMware Technical Journal*, Vol. 1(1), April 2012.
- P. Xiong et al. "vPerfGuard: An automated model-driven framework for application performance diagnosis in consolidated cloud environments." *ICPE 2013*.
- A. Gulati, "Towards proactive resource management in virtualized datacenters." *RESOLVE 2013*.
- L. Lu, et al., "Application-Driven dynamic vertical scaling of virtual machines in resource pools." to appear at *NOMS 2014*.



vmware

vmware



## MODERN INFRASTRUCTURE: THE CONVERGENCE OF NETWORK, COMPUTE, AND DATA

**Jason Hoffman, Ericsson**

The three pillars of our industry are network, compute, and data. All trends come down to the convergence of these. The convergence of network and compute resulted in the “the network is the computer”; the convergence of network and data spawned the entire networked storage industry and now we believe we’re in the technology push where we are converging compute and data. In this talk, we’ll cover the philosophical basis, the overall architecture, and the deep details of a holistic datacenter implementation.



# SOME THOUGHTS

Johan Eker  
Principal Researcher  
Ericsson

Some slides and ideas courtesy of Jason Holman, Head of the Ericsson Cloud System and Platforms

## OPERATOR CLOUD SEGMENTS

<b>Operator Public Cloud</b>	<b>Operator IT Private Cloud</b>	<b>Operator Telecom Private Cloud</b>
Virtualized Compute + WAN resources sold as a service to Enterprises aka "Managed Cloud"	Virtualized IT functions OSS, BSS, SDP, ERP, CRM etc	Virtualized telecom functions (e.g. IMS)

## ERICSSON, WHO?



NFV  
Network Functions  
Virtualization

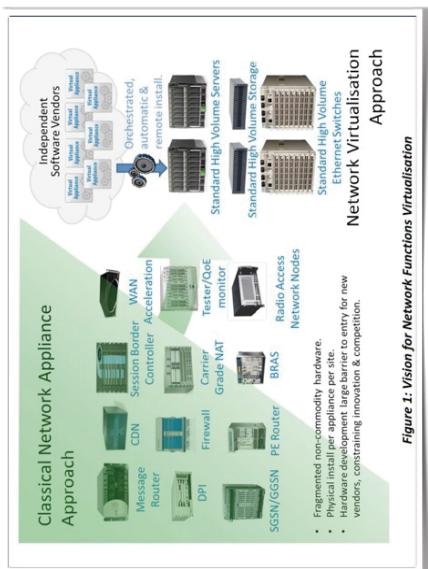


NFV

Network  
Functions

Virtualization





## 500 BILLION DEVICES

## 500 BILLION DEVICES

Collecting data. Acting on data.

## 500 BILLION DEVICES

Doing what?

## DATA.

Source. Observer. Student.

## DATA FROM HUMANS

Photos. Videos. Text. Audio.



## DATA FROM HUMANS

Google. FB. Enterprise File & Mail.

## DATA FROM MACHINES

Servers, Phones, Wearables, Sensors, Cars.

## DATA FROM NATURE

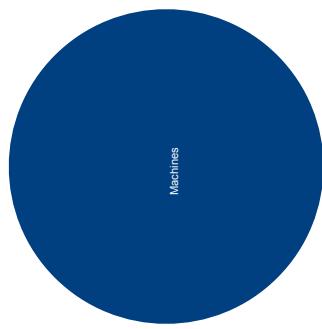
Is the highest resolution data available.

10,000,000 GENOMES IS 20EB  
Keep all the data

A SINGLE IDEA CAN CONSUME  
AN ENTIRE INDUSTRY

STORAGE INDUSTRY SHIPPED 16EB  
In 2012 (according to IDC)

## WHAT WE'VE BEEN DOING



Machines

Humans

Humans Source, Humans Observe, Humans Learn  
Nature Source, Nature Observe, Nature Learns  
Nature Source, Humans Observe, Humans Learn  
Nature Source, Machine Observes, Humans Learn



## WHAT'S NEW



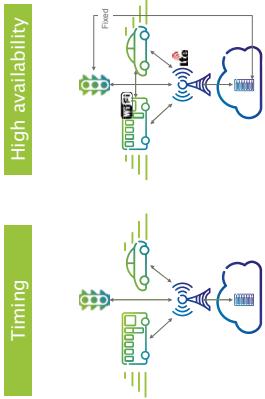
Nature

500 BILLION DEVICES  
Collecting Data.  
Learning.  
Acting.

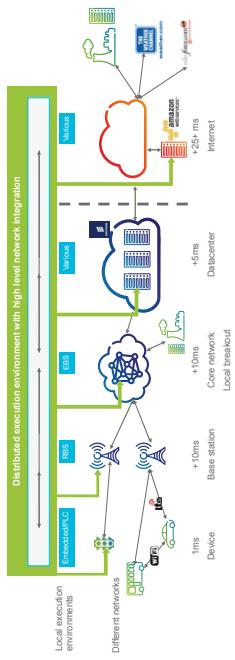


Humans Source, Machines Observe, Humans Learn  
Humans Source, Machines Observe, Machines Learn  
Machines Source, Machines Observe, Machines Learn  
Nature Source, Machines Observe, Machines Learn

## MISSION CRITICAL CLOUD



## EVERYTHING IS PROGRAMMABLE



## THIS IS THE HARDWARE



## HOW TO PROGRAM?

Distributed. Heterogenous. Dynamic.

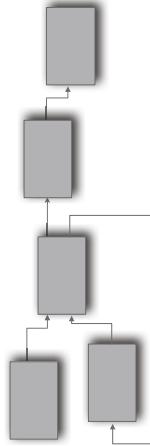
## HOW TO PROGRAM?

Ease of use. Resource. Timing. Resilient.

Distributed. Network, Computed & Data integration

## AN APPLICATION AS A GRAPH

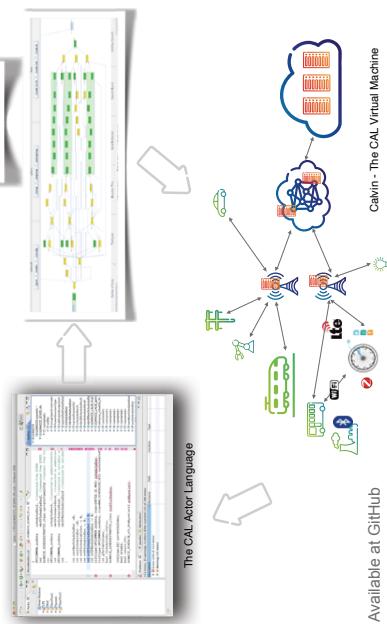
- Nodes called actors
- Message passing
- Actors only interact via ports & FIFO connections
- Scheduling decoupled from algorithm
- Programming in CAL, Erlang or X



## HOW TO MANAGE?

Distributed. Network, Computed & Data integration

## CALTOOPIA.ORG



Available at GitHub

Calvin - The CAL Virtual Machine

## THE END



Everything is programmable

The data collected is used to control things

The cloud is integrated with devices & network

**EVENT-BASED CONTROL: A WAY TO REDUCE  
RECONFIGURATION IN AUTONOMIC COMPUTING**  
**Nicholas Marchand, GIPSA lab, France**

My talk will focus on a new control techniques called event-based control. This approach recently developed differs from classical control in the sense that the control value is updated only when needed. Reducing the number of control update often means for systems in the computer science domain a reduction of system re-configuration. The talk will focus on the practical use of this new technique to control the service time of an Hadoop MapReduce cluster. A comparison between classical control and event-based control will be given.



**Outline**



**Event based control:**  
A way to reduce reconfigurations in autonomic computing ?

N. Marchand  
gipsa, Control Systems Department, Grenoble, FRANCE

**Motivation**

- Dealing with the dynamics: time is crucial
- Mathematical tools to "control" a system
- By "control", we mean being able to
  - define a control objective
  - define control actions accordingly
  - guarantee performances of the controlled system
  - despite errors
  - despite perturbations
  - Facing everything that is unknown
  - Guarantee stability
- Many other area of control theory are relevant to computer science
- Fault tolerant control, fault detection, supervision, etc.
- Nowadays control theory is everywhere...
- automotive, robotics energy (grids, production, etc.), microelectronics (DVFS), etc.
- ...except maybe in computer science

**Introduction**

- Motivation
- Outline
- Outline of the talk

**A highly nonlinear example**

**Event-based PID controller**

- Formulation
- Cases
- Simulations
- MapReduce control
- EB-Controller
- Conclusion

**Conclusion**

**Challenging difficulties**

Let's start with the hardest things

LCCC workshop on Cloud Control	N. Marchand (eipss)	gipsa-lab	LCCC workshop on Cloud Control	N. Marchand (eipss)	gipsa-lab	LCCC workshop on Cloud Control	N. Marchand (eipss)	gipsa-lab
22/03/2012 1 / 20	22/03/2012 1 / 20	22/03/2012 1 / 20	22/03/2012 2 / 20	22/03/2012 2 / 20	22/03/2012 2 / 20	22/03/2012 3 / 20	22/03/2012 3 / 20	22/03/2012 4 / 20



**Control theory**

Languages difficulties	Words	Computer science	Control theory
Autonomic or autonomous	Autonomic	Controlled	Uncontrolled
Time response	Time	...	time needed to reach x% of the final value
Parameter	Parameter	...	constants
Cloud	Cloud	Set of interconnected computers	Look at Lund's sky
Control	Control	Parametrization	$\frac{dx}{dt} = f(x,u)$
...	...	...	...

**Interest of both communities**

- No physics behind algorithms, applications, services, etc.
- Let's do things in cloud !! (Sara Bouchenek from LiG-lab)



## dip-sa-lab

### Challenging difficulties

Let's start with the hardest things

## dip-sa-lab

### Challenging difficulties

Let's start with the hardest things

Languages difficulties	
Words	Computer science
Autonomic or Autonomous	Controlled
Time response	Uncontrolled
Parameter	Queuing + processing time variables you can change
Cloud	Set of interconnected computers
Control	Parametrization
...	...
No phys	
! Let's do	Control for control theorist

Languages difficulties	
Words	Control theory
Autonomic or Autonomous	Uncontrolled
Time response	time needed to reach x% of the final value constants
Parameter	Look at Lund's sky
Cloud	$\frac{dx}{dt} = f(x, u)$
Control	...
...	...
No phys	
! Let's do	Control for control theorist

Languages difficulties	
Words	Control theory
Autonomic or Autonomous	Uncontrolled
Time response	time needed to reach x% of the final value constants
Parameter	Look at Lund's sky
Cloud	$\frac{dx}{dt} = f(x, u)$
Control	...
...	...
No phys	
! Let's do	Control for control theorist

Languages difficulties	
Words	Control theory
Autonomic or Autonomous	Uncontrolled
Time response	time needed to reach x% of the final value constants
Parameter	Look at Lund's sky
Cloud	$\frac{dx}{dt} = f(x, u)$
Control	...
...	...
No phys	
! Let's do	Control for control theorist

Languages difficulties	
Words	Control theory
Autonomic or Autonomous	Uncontrolled
Time response	time needed to reach x% of the final value constants
Parameter	Look at Lund's sky
Cloud	$\frac{dx}{dt} = f(x, u)$
Control	...
...	...
No phys	
! Let's do	Control for control theorist

Languages difficulties	
Words	Control theory
Autonomic or Autonomous	Uncontrolled
Time response	time needed to reach x% of the final value constants
Parameter	Look at Lund's sky
Cloud	$\frac{dx}{dt} = f(x, u)$
Control	...
...	...
No phys	
! Let's do	Control for control theorist

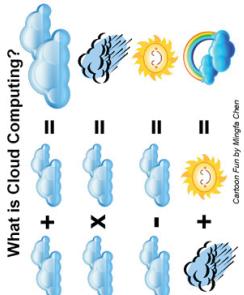
Languages difficulties	
Words	Control theory
Autonomic or Autonomous	Uncontrolled
Time response	time needed to reach x% of the final value constants
Parameter	Look at Lund's sky
Cloud	$\frac{dx}{dt} = f(x, u)$
Control	...
...	...
No phys	
! Let's do	Control for control theorist

Languages difficulties	
Words	Control theory
Autonomic or Autonomous	Uncontrolled
Time response	time needed to reach x% of the final value constants
Parameter	Look at Lund's sky
Cloud	$\frac{dx}{dt} = f(x, u)$
Control	...
...	...
No phys	
! Let's do	Control for control theorist

## Challenging difficulties

Let's start with the hardest things

- Languages difficulties
- Interest of both communities
- No physics behind algorithms, applications, services, etc.
- "Let's do things in cloud computing" (Sara Bouchenak from LIG-lab)



## Menu

- Starter: An short example of how bad computer system can be for control theory:

- Admission control for PostgreSQL database server

### Main dish:

- Cloud control needs to
  - react to spikes (high frequency)
  - reconfigure as less as possible (low frequency)
  - Antinomic !
- Focus on Event-Based control
  - More on event-based PID
  - Short presentation of extensions
  - Assure SLA compliance in Hadoop Mapreduce
- Dessert: What need to be more efficient

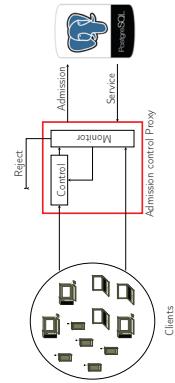
Hope it will be not too indigestible !

## Challenging difficulties

Let's start with the hardest things

- Languages difficulties
- Interest of both communities
- No physics behind algorithms, applications, services, etc.
- "Let's do things in cloud control" (Sara Bouchenak from LIG-lab)

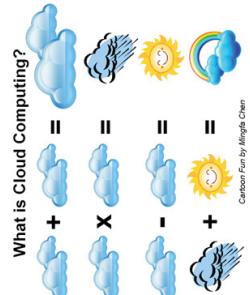
## A nonlinear system example



## gipsa-lab

gipsa-lab

- Languages difficulties
- Interest of both communities
- No physics behind algorithms, applications, services, etc.
- "Let's do things in cloud computing" (Sara Bouchenak from LIG-lab)



## gipsa-lab

gipsa-lab

- Starter: An short example of how bad computer system can be for control theory:

- Admission control for PostgreSQL database server

### Main dish:

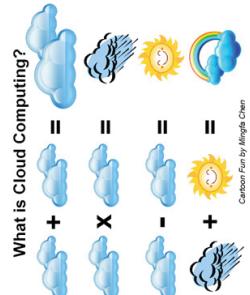
- Cloud control needs to
  - react to spikes (high frequency)
  - reconfigure as less as possible (low frequency)
  - Antinomic !
- Focus on Event-Based control
  - More on event-based PID
  - Short presentation of extensions
  - Assure SLA compliance in Hadoop Mapreduce
- Dessert: What need to be more efficient

Hope it will be not too indigestible !

## gipsa-lab

gipsa-lab

- Languages difficulties
- Interest of both communities
- No physics behind algorithms, applications, services, etc.
- "Let's do things in cloud computing" (Sara Bouchenak from LIG-lab)



## gipsa-lab

gipsa-lab

- Starter: An short example of how bad computer system can be for control theory:

- Admission control for PostgreSQL database server

### Main dish:

- Cloud control needs to
  - react to spikes (high frequency)
  - reconfigure as less as possible (low frequency)
  - Antinomic !
- Focus on Event-Based control
  - More on event-based PID
  - Short presentation of extensions
  - Assure SLA compliance in Hadoop Mapreduce
- Dessert: What need to be more efficient

Hope it will be not too indigestible !

## N. Marchand (eipsa)

N. Marchand

## LCCC workshop on Cloud Control

22/03/2012

## N. Marchand (eipsa)

N. Marchand

## LCCC workshop on Cloud Control

22/03/2012

## LCCC workshop on Cloud Control

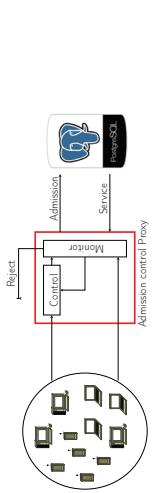
22/03/2012

## LCCC workshop on Cloud Control

22/03/2012

## A nonlinear system example

gipsa-lab



- Model gives when saturated:
 
$$\begin{cases} \frac{dN}{dt} = (1-\alpha) \cdot T_i - T_o \\ \frac{d\alpha}{dt} = -\frac{1}{\Delta} \left[ \alpha - \frac{N}{MPL} (1 - \frac{T_o}{T_i}) \right] \\ \frac{dT_o}{dt} = -\frac{1}{\Delta} \left[ T_o - \frac{N}{aN^2 + bN + c} \right] \end{cases}$$

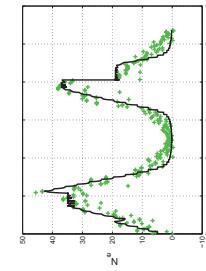
## A nonlinear system example

gipsa-lab

Model gives when saturated:

$$\begin{cases} \frac{dN}{dt} = (1-\alpha) \cdot T_i - T_o \\ \frac{d\alpha}{dt} = -\frac{1}{\Delta} \left[ \alpha - \frac{N}{MPL} (1 - \frac{T_o}{T_i}) \right] \\ \frac{dT_o}{dt} = -\frac{1}{\Delta} \left[ T_o - \frac{N}{aN^2 + bN + c} \right] \end{cases}$$

- Quite a pretty model
  - few variables/parameters
  - easily identifiable
  - fits well



LCCC workshop on Cloud Control  
22/03/2012 7 / 20

## A nonlinear system example

gipsa-lab

Model gives when saturated:

$$\begin{cases} \frac{dN}{dt} = (1-\alpha) \cdot T_i - T_o \\ \frac{d\alpha}{dt} = -\frac{1}{\Delta} \left[ \alpha - \frac{N}{MPL} (1 - \frac{T_o}{T_i}) \right] \\ \frac{dT_o}{dt} = -\frac{1}{\Delta} \left[ T_o - \frac{N}{aN^2 + bN + c} \right] \end{cases}$$

- Quite a pretty model
  - few variables/parameters
  - easily identifiable
  - fits well

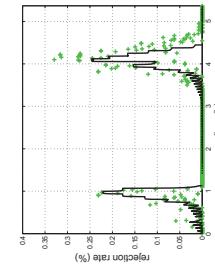
## A nonlinear system example

gipsa-lab

Model gives when saturated:

$$\begin{cases} \frac{dN}{dt} = (1-\alpha) \cdot T_i - T_o \\ \frac{d\alpha}{dt} = -\frac{1}{\Delta} \left[ \alpha - \frac{N}{MPL} (1 - \frac{T_o}{T_i}) \right] \\ \frac{dT_o}{dt} = -\frac{1}{\Delta} \left[ T_o - \frac{N}{aN^2 + bN + c} \right] \end{cases}$$

- Quite a pretty model
  - few variables/parameters
  - easily identifiable
  - fits well



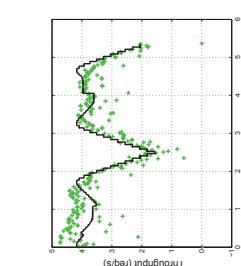
LCCC workshop on Cloud Control  
22/03/2012 7 / 20

## A nonlinear system example



- Model gives when saturated:

$$\begin{cases} \frac{dN}{dt} = (1-\alpha) \cdot T_i - T_o \\ \frac{d\alpha}{dt} = -\frac{1}{\Delta} \left[ \alpha - \frac{N}{MPL} (1 - \frac{T_o}{T_i}) \right] \\ \frac{dT_o}{dt} = -\frac{1}{\Delta} \left[ T_o - \frac{N}{aN^2 + bN + c} \right] \end{cases}$$



- Quite a pretty model

- few variables/parameters
- easily identifiable
- fits well

- Model gives when saturated:

✓ Quite a pretty model

✗ Not an easy model

- Model is highly nonlinear
- Not in the standard form for control theory
- Control is the level saturation of an exogenous input

## A nonlinear system example



- Controlled thanks nonlinear control theory (Lyapunov)
- Stability is guaranteed for any value of the parameters
- No danger of miss-identification

- Controlled thanks nonlinear control theory (Lyapunov)
- Stability is guaranteed for any value of the parameters
- No danger of miss-identification

Introduction  
Motivation  
Outline  
**NL example**

## A nonlinear system example

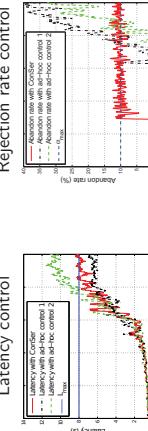


- Controlled thanks nonlinear control theory (Lyapunov)
- Stability is guaranteed for any value of the parameters
- No danger of miss-identification

- Controlled thanks nonlinear control theory (Lyapunov)
- Stability is guaranteed for any value of the parameters
- No danger of miss-identification

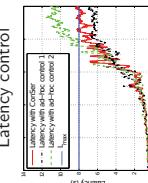
Introduction  
Motivation  
Outline  
**NL example**

## Rejection rate control



22/03/2012 7 / 26

## latency control



22/03/2012 7 / 26

## A nonlinear system example



- Model gives when saturated:

$$\begin{cases} \frac{dN}{dt} = (1-\alpha) \cdot T_i - T_o \\ \frac{d\alpha}{dt} = -\frac{1}{\Delta} \left[ \alpha - \frac{N}{MPL} (1 - \frac{T_o}{T_i}) \right] \\ \frac{dT_o}{dt} = -\frac{1}{\Delta} \left[ T_o - \frac{N}{aN^2 + bN + c} \right] \end{cases}$$



- Quite a pretty model

- few variables/parameters
- easily identifiable
- fits well

- Model gives when saturated:

✓ Quite a pretty model

✗ Not an easy model

- Model is highly nonlinear
- Not in the standard form for control theory
- Control is the level saturation of an exogenous input

## A nonlinear system example

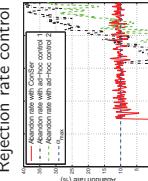


- Model gives when saturated:

- Model gives when saturated:

- Model is highly nonlinear
- Not in the standard form for control theory
- Control is the level saturation of an exogenous input

## Rejection rate control



22/03/2012 8 / 26

## latency control



22/03/2012 8 / 26

## Menu



### Event-based sampling vs. periodic sampling

- LCCC workshop on Cloud Control N. Marchand**
- Introduction Motivation Outlines
- NL example**
- EB-PID Formulation Cases Simulations MapReduce control EB-Control Conclusion
- Main dish:**
  - Cloud control needs to
    - react to spikes (high frequency)
    - reconfigure as less as possible (low frequency)
    - Antinomic !
  - Focus on Event-Based control
  - More on event based PID
  - Short presentation of extensions
  - Assure SLA compliance in Hadoop Mapreduce



N. Marchand (épia)

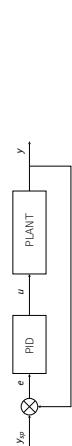
gipsa-lab

LCCC workshop on Cloud Control N. Marchand (épia) 22/03/2012 9 / 26 N. Marchand (épia) 22/03/2012 10 / 26



### PID controller

#### Classical work



#### Classical PID

- In the frequency domain:

$$U(s) = K \left( E(s) + \frac{1}{T_d s} E(s) + T_d s E(s) \right)$$

- Discrete time version ( $h_{nom}$ : sampling period,  $N$  tunes the filter):

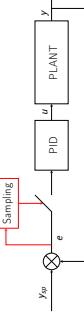
$$\begin{aligned} u_p(t_k) &= K e(t_k) \\ u_l(t_{k+1}) &= u_l(t_k) + K_l h_{nom} e(t_k) \\ u_d(t_k) &= \frac{K e(t_k)}{T_d + N h_{nom}} u_d(t_{k-1}) + \frac{K T_d N}{T_d + N h_{nom}} (e(t_k) - e(t_{k-1})) \\ u &= u_p + u_l + u_d \end{aligned}$$

- LCCC workshop on Cloud Control N. Marchand**
- Introduction Motivation Outlines
- NL example**
- EB-PID Formulation Cases Simulations MapReduce control EB-Control Conclusion
- Periodic sampling**
  - Sampling periodically on time
  - Analogical to Riemann's integral
  - Well known theory (Shannon, etc.)
- Event-based sampling**
  - Sampling on level's
  - At first glance close to Lebesgues integral
  - Different extension :
  - Outside event
  - State/output dependent sampling (self-triggered)
  - Should reduce transmission/computation
  - Few theory

LCCC workshop on Cloud Control N. Marchand (épia) 22/03/2012 11 / 20 N. Marchand (épia) 22/03/2012 12 / 20

### PID controller

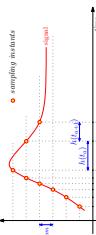
#### EB-PID



#### Idea:

- do not update the control if  $y$  is close to  $y_{sp}$ , typically if  $\|e(t_a) - e(t_{a-1})\| \leq q_{nom}$

- No need to respect Shannon
- Bad behavior of the integral part
- $q_{nom}$  linked to precision and noise



LCCC workshop on Cloud Control N. Marchand (épia) 22/03/2012 12 / 20 N. Marchand (épia) 22/03/2012 13 / 20

## What can happen (1/3) ?

- Simple integrator  $\frac{dx}{dt} = u$
- Level-crossing sampling  $\Rightarrow$  update the control when  $e(x) = 0$
- Trajectory :  $x_{i+1} = x_i + (t_{i+1} - t_i) \cdot u$
- Sampling instants**  $t_i$
- Sampling set:**  $T_{e,k,x_0} := \{t_i\} \supset \{0\}$
- NL example**
- EB-PID**
- Formulation**
- Cases**
- Simulations**
- Motivation**
- Outline**
- Conclusion**

- ④  $k(x) = -x$ ,  $e(x) = 0$  when  $|x| = \exp(-\kappa)$ ,  $\kappa \in \mathbb{Z}$
- $T_{e,k,x_0} := \{j(1 - \exp(-1)) / j \in \mathbb{N}\}$
  - closed-loop system is globally asymptotically stable
  - the solution is defined on  $[0, \infty]$
  - the sampling set depends upon  $x_0$
  - $k(x) = -x^{\frac{1}{\kappa}}$ ,  $e(x) = 0$  when  $|x| = \frac{1}{\kappa}$ ,  $\kappa \in \mathbb{Z}$
  - $x_0 = 1 \Rightarrow t_{i+1} - t_i = \frac{1}{\sqrt{i(i+1)}}$
  - closed-loop system is globally asymptotically stable
  - Zero phenomenon: the solution is defined only on  $[0, 1.86[$

## What can happen (1/3) ?

- Simple integrator  $\frac{dx}{dt} = u$
- Level-crossing sampling  $\Rightarrow$  update the control when  $e(x) = 0$
- Trajectory :  $x_{i+1} = x_i + (t_{i+1} - t_i) \cdot u$
- Sampling instants**  $t_i$
- Sampling set:**  $T_{e,k,x_0} := \{t_i\} \supset \{0\}$
- NL example**
- EB-PID**
- Formulation**
- Cases**
- Simulations**
- Motivation**
- Outline**
- Conclusion**

- ④  $k(x) = -x$ ,  $e(x) = 0$  when  $|x| = \exp(-\kappa)$ ,  $\kappa \in \mathbb{Z}$
- $T_{e,k,x_0} := \{j \cdot (1 - \exp(-1)) / j \in \mathbb{N}\}$
  - closed-loop system is globally asymptotically stable
  - the solution is defined on  $[0, \infty]$
  - the sampling set depends upon  $x_0$
  - $k(x) = -x^{\frac{1}{\kappa}}$ ,  $e(x) = 0$  when  $|x| = \frac{1}{\kappa}$ ,  $\kappa \in \mathbb{Z}$
  - $x_0 = 1 \Rightarrow t_{i+1} - t_i = \frac{1}{\sqrt{i(i+1)}}$
  - closed-loop system is globally asymptotically stable
  - Zero phenomenon: the solution is defined only on  $[0, 1.86[$



- LCCC workshop on Cloud Control  
N. Marchand
- Introduction  
Motivation  
Outline  
Conclusion

- LCCC workshop on Cloud Control  
N. Marchand (gipsa)
- Introduction  
Motivation  
Outline  
Conclusion

## What can happen (1/3) ?

- Simple integrator  $\frac{dx}{dt} = u$
- Level-crossing sampling  $\Rightarrow$  update the control when  $e(x) = 0$
- Trajectory :  $x_{i+1} = x_i + (t_{i+1} - t_i) \cdot u$
- Sampling instants**  $t_i$
- Sampling set:**  $T_{e,k,x_0} := \{t_i\} \supset \{0\}$
- NL example**
- EB-PID**
- Formulation**
- Cases**
- Simulations**
- Motivation**
- Outline**
- Conclusion**

- ④  $k(x) = -x^{\frac{1}{\kappa}}$ ,  $e(x) = 0$  when  $|x| = \frac{1}{\kappa}$ ,  $\kappa \in \mathbb{Z}$
- $x_0 = 1 \Rightarrow t_{i+1} - t_i = \frac{1}{\sqrt{i(i+1)}}$
  - closed-loop system is globally asymptotically stable
  - Zero phenomenon: the solution is defined only on  $[0, 1.86[$



- LCCC workshop on Cloud Control  
N. Marchand

- LCCC workshop on Cloud Control  
N. Marchand (gipsa)
- Introduction  
Motivation  
Outline  
Conclusion

## What can happen (2/3) ?

- Simple integrator  $\frac{dx}{dt} = u$
- Level-crossing sampling  $\Rightarrow$  update the control when  $e(x) = 0$
- Trajectory :  $x_{i+1} = x_i + (t_{i+1} - t_i) \cdot u$
- Sampling instants**  $t_i$
- Sampling set:**  $T_{e,k,x_0} := \{t_i\} \supset \{0\}$
- NL example**
- EB-PID**
- Formulation**
- Cases**
- Simulations**
- Motivation**
- Outline**
- Conclusion**

- ④  $k(x) = -x^{\frac{1}{\kappa}}$ ,  $e(x) = 0$  when  $|x| = \frac{1}{\kappa}$ ,  $\kappa \in \mathbb{Z}$
- $x_0 = 1 \Rightarrow t_{i+1} - t_i = \frac{1}{\sqrt{i(i+1)}}$
  - closed-loop system is globally asymptotically stable
  - $\lim t_{i+1} - t_i = 0$  when  $\lim t_i = \infty$
  - Infinitely fast sampling at infinity
  - $k(x) = -x^{\frac{1}{\kappa}}$ ,  $e(x) = 0$  when  $|x| = \exp(-\kappa)$ ,  $\kappa \in \mathbb{Z}$
  - $x_0 = 1 \Rightarrow t_{i+1} - t_i = \exp(2/i) \cdot [1 - \exp(-1)]$
  - closed-loop system is globally asymptotically stable
  - $\lim t_{i+1} - t_i = \infty$  when  $\lim t_i = \infty$
  - Shannon's condition is inconsistent
  - the solution is defined on  $[0, \infty]$
  - the sampling set depends upon  $x_0$
  - Infinitely slow sampling at infinity



- LCCC workshop on Cloud Control  
N. Marchand (gipsa)
- Introduction  
Motivation  
Outline  
Conclusion



- LCCC workshop on Cloud Control  
N. Marchand (gipsa)
- Introduction  
Motivation  
Outline  
Conclusion

## What can happen (2/3) ?



- ③  $k(x) = -x$ ,  $e(x) = 0$  when  $|x| = \frac{1}{k}$ ,  $x \in \mathbb{Z}$
- $x_0 = 1 \Rightarrow t_{i+1} - t_i = \frac{1}{\sqrt{k}}$

- closed-loop system is **globally asymptotically stable**
- $\lim t_{i+1} - t_i = 0$  when  $\lim t_i = \infty$
- Infinitely fast sampling at infinity

- ④  $k(x) = -x^3$ ,  $e(x) = 0$  when  $|x| = \exp(-x)$ ,  $x \in \mathbb{Z}$
- $x_0 = 1 \Rightarrow t_{i+1} - t_i = [\exp(2)] \cdot [1 - \exp(-1)]$

- closed-loop system is **globally asymptotically stable**
- $\lim t_{i+1} - t_i = \infty$  when  $\lim t_i = \infty$
- Shannon's condition is inconsistent
- the solution is defined on  $[0, \infty]$
- Infinitely slow sampling at infinity

## What can happen (3/3) ?



- **Unstable system:**  $\frac{dx}{dt} = (x+u)^3$
- Solution is:  $x_{i+1} = \frac{x_i+u}{\sqrt{1-2(t_{i+1}-t_i)(x_i+u)^2}} - u$

**N. Marchand**

Introduction  
Motivation  
Outline  
NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

Introduction  
Motivation  
Outline

NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

Introduction  
Motivation  
Outline  
NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

Introduction  
Motivation  
Outline  
NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

## What can happen (3/3) ?



- **Unstable system:**  $\frac{dx}{dt} = (x+u)^3$
- Solution is:  $x_{i+1} = \frac{x_i+u}{\sqrt{1-2(t_{i+1}-t_i)(x_i+u)^2}} - u$

**N. Marchand**

Introduction  
Motivation  
Outline  
NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

Introduction  
Motivation  
Outline

NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

Introduction  
Motivation  
Outline  
NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

- $k(x) = -2x$ ,  $e(x) = 0$  when  $|x| = \exp(-x)$ ,  $x \in \mathbb{Z}$  and initial condition  $x_0 = 1$
- $t_{i+1} - t_i = \frac{\exp(2)}{2} \cdot \left[ 1 - \frac{1}{(2-\exp(-1))^2} \right]$

- closed-loop system is **globally asymptotically stable**
- $\lim t_{i+1} - t_i = \infty$  when  $\lim t_i = \infty$
- Shannon's condition is inconsistent
- the solution is defined on  $[0, \infty]$

$$u_i(t_a) = u_i(t_{a-1}) + K_i h e(t_a)$$

limited

- Saturation, Exponential forgetting factor, Hybrid, etc.

## What can happen (3/3) ?



- **Unstable system:**  $\frac{dx}{dt} = (x+u)^3$
- Solution is:  $x_{i+1} = \frac{x_i+u}{\sqrt{1-2(t_{i+1}-t_i)(x_i+u)^2}} - u$

**N. Marchand**

Introduction  
Motivation  
Outline  
NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

Introduction  
Motivation  
Outline

NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

Introduction  
Motivation  
Outline  
NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

## PID controller



- We focus on the integral part
- What happens one waits too long before updating the control?

**N. Marchand**

Introduction  
Motivation  
Outline  
NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

Introduction  
Motivation  
Outline  
NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

Introduction  
Motivation  
Outline  
NL example  
EB-PID  
Formulation  
**Case**  
Simulation  
MapReduce control  
EB-Control  
Conclusion

- Strong overshoot when the control is updated (similar to saturated PID without antwindup)
- Solution: replace the product  $h \cdot e$  by a bounded function  $he$ :

$$u_i(t_a) = u_i(t_{a-1}) + K_i he(t_a)$$

limited

22/03/2012 15 / 20

22/03/2012 15 / 20

22/03/2012 15 / 20

22/03/2012 15 / 20

22/03/2012 15 / 20

22/03/2012 15 / 20

22/03/2012 15 / 20

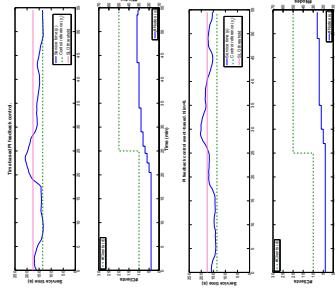
22/03/2012 15 / 20



## MapReduce control

Simple PI controller

- Time based (8 updates) vs. Event based (4 updates)



## More food for people in control theory

- Event-based control is really recent
- Now exist :
  - Almost basic linear control where carried in an event-based framework (PID, LQR, etc.)
  - Sontag's general formula has been extended
  - A lot of strategies based on Lyapunov theory exist
  - Many practical implementations even on noisy and unstable systems
  - Early results are appearing for time-delayed systems
- Remain to clarify
  - Real number of control updates
  - All what it brings (in good and bad) is not clear
  - Frequency analysis is less convenient

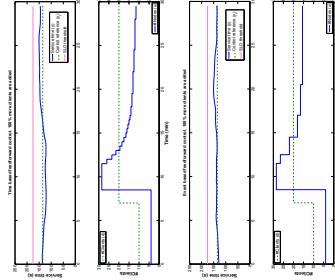
N. Marchand (gipsa)



## MapReduce control

Simple PI controller with feedforward

- Time based (18 updates) vs. Event based (6 updates)



## Conclusion

- People always adopt control theory ...
  - Ecological constraints: car industry (in the 90's)
  - Nuclear plant: from the beginning (and one must be sure it works)
  - Cost constraints: Petrol industries (in late 50's)
  - Energy constraints: Embedded systems (in the 00's)
  - Crash risks: Smart Grids (nowadays)
- In all cases it was (is) a question of money
  - Is it a question of money in cloud computing ?
  - Before adopting control theory, intuitive control was the strategy
  - Theory is the only way (control theory, game theory, queuing theory etc.)
- to face safety complexity
- to guarantee results (even in unknown/unpredictable environment)
- to have flexibility

N. Marchand (gipsa)



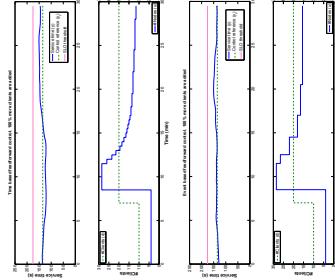
N. Marchand (gipsa)



## MapReduce control

EB-control

- Time based (18 updates) vs. Event based (6 updates)



## Conclusion

- People always adopt control theory ...
  - Ecological constraints: car industry (in the 90's)
  - Nuclear plant: from the beginning (and one must be sure it works)
  - Cost constraints: Petrol industries (in late 50's)
  - Energy constraints: Embedded systems (in the 00's)
  - Crash risks: Smart Grids (nowadays)
- In all cases it was (is) a question of money
  - Is it a question of money in cloud computing ?
  - Before adopting control theory, intuitive control was the strategy
  - Theory is the only way (control theory, game theory, queuing theory etc.)
- to face safety complexity
- to guarantee results (even in unknown/unpredictable environment)
- to have flexibility

N. Marchand (gipsa)



N. Marchand (gipsa)





## Conclusion

### • People always adopt control theory ... under constraint

- Ecological constraints: car industry (in the 90's)
- Nuclear plant: from the beginning (and one must be sure it works)
- Cost constraints: Petrol industries (in late 50's)
- Energy constraints: Embedded systems (in the 00's)
- Crash risks: Smart Grids (nowadays)
- In all cases it was (is) a question of money
- Is it a question of money in cloud computing ?
- Before adopting control theory, intuitive control was the strategy
- Theory is the only way (control theory, game theory, queuing theory, etc.)
- to face safely complexity
- to guarantee results (even in unknown/unpredictable environment)
- to have flexibility



## Conclusion

### • People always adopt control theory ... under constraint

- Ecological constraints: car industry (in the 90's)
- Nuclear plant: from the beginning (and one must be sure it works)
- Cost constraints: Petrol industries (in late 50's)
- Energy constraints: Embedded systems (in the 00's)
- Crash risks: Smart Grids (nowadays)
- In all cases it was (is) a question of money
- Is it a question of money in cloud computing ?
- Before adopting control theory, intuitive control was the strategy
- Theory is the only way (control theory, game theory, queuing theory, etc.)
- to face safely complexity
- to guarantee results (even in unknown/unpredictable environment)
- to have flexibility



## Conclusion

### • People always adopt control theory ... under constraint

- Ecological constraints: car industry (in the 90's)
- Nuclear plant: from the beginning (and one must be sure it works)
- Cost constraints: Petrol industries (in late 50's)
- Energy constraints: Embedded systems (in the 00's)
- Crash risks: Smart Grids (nowadays)
- In all cases it was (is) a question of money
- Is it a question of money in cloud computing ?
- Before adopting control theory, intuitive control was the strategy
- Theory is the only way (control theory, game theory, queuing theory, etc.)
- to face safely complexity
- to guarantee results (even in unknown/unpredictable environment)
- to have flexibility



## Conclusion

### • People always adopt control theory ... under constraint

- Ecological constraints: car industry (in the 90's)
- Nuclear plant: from the beginning (and one must be sure it works)
- Cost constraints: Petrol industries (in late 50's)
- Energy constraints: Embedded systems (in the 00's)
- Crash risks: Smart Grids (nowadays)
- In all cases it was (is) a question of money
- Is it a question of money in cloud computing ?
- Before adopting control theory, intuitive control was the strategy
- Theory is the only way (control theory, game theory, queuing theory, etc.)
- to face safely complexity
- to guarantee results (even in unknown/unpredictable environment)
- to have flexibility



## Conclusion

- Ecological constraints: car industry (in the 90's)
- Nuclear plant: from the beginning (and one must be sure it works)
- Cost constraints: Petrol industries (in late 50's)
- Energy constraints: Embedded systems (in the 00's)
- Crash risks: Smart Grids (nowadays)
- In all cases it was (is) a question of money
- Is it a question of money in cloud computing ?
- Before adopting control theory, intuitive control was the strategy
- Theory is the only way (control theory, game theory, queuing theory, etc.)
- to face safely complexity
- to guarantee results (even in unknown/unpredictable environment)
- to have flexibility



## Conclusion



- People always adopt control theory ... **under constraint**
  - Ecological constraints: car industry (in the 90's)
  - Nuclear plant: from the beginning (and one must be sure it works)
  - Cost constraints: Petrol industries (in late 50's)
  - Energy constraints: Embedded systems (in the 00's)
  - Crash risks: Smart Grids (nowadays)
- In all cases it was (is) a question of money
- Is it a question of money in cloud computing ?
- Before adopting control theory, intuitive control was the strategy
- Theory is the only way (control theory, game theory, queuing theory, etc.)
- to face **safely** complexity
- to guarantee results (even in unknown/unpredictable environment)
- to have flexibility

## What need to be improved



- From the computer science side:
  - Classification of problems in big classes
  - Standardisation of inputs/outputs/variables for each class
  - Co-design / Control aware software
  - Better sort things by speed
  - Patience (to explain and to get results)
- From the control theory side:
  - More interest
  - Building a theory that handles computer science problems
- From both side:
  - Spend more time together
  - Mix techniques from both sides
- Some inspiring fields
  - Embedded systems
  - deadline problems, energy optimization, re-allocation, heterogeneous MPSoC, ...



## What need to be improved



- From the computer science side:
  - Classification of problems in big classes
  - Standardisation of inputs/outputs/variables for each class
  - Co-design / Control aware software
  - Better sort things by speed
  - Patience (to explain and to get results)
- From the control theory side:
  - More interest
  - Building a theory that handles computer science problems
- From both side:
  - Spend more time together
  - Mix techniques from both sides
- Some inspiring fields
  - Embedded systems
  - deadline problems, energy optimization, re-allocation, heterogeneous MPSoC, ...

## What need to be improved



- From the computer science side:
  - Classification of problems in big classes
  - Standardisation of inputs/outputs/variables for each class
  - Co-design / Control aware software
  - Better sort things by speed
  - Patience (to explain and to get results)
- From the control theory side:
  - More interest
  - Building a theory that handles computer science problems
- From both side:
  - Spend more time together
  - Mix techniques from both sides
- Some inspiring fields
  - Embedded systems
  - deadline problems, energy optimization, re-allocation, heterogeneous MPSoC, ...
- Electrical grids
  - centralized/decentralized, providers/consumers, cascading failure, load management, etc.



## What need to be improved

- From the computer science side:
  - Classification of problems in big classes
  - Standardisation of inputs/outputs/variables for each class
  - Co-design / Control aware software
  - Better sort things by speed
  - Patience (to explain and to get results)
- From the control theory side:
  - More interest
  - Building a theory that handles computer science problems
- From both side:
  - Spend more time together
  - Mix techniques from both side
- Some inspiring fields
  - Embedded systems
  - deadline problems, energy optimization, re-allocation, heterogeneous MPSoC, ...
  - Electrical grids
  - centralized/decentralized, providers/consumers, cascading failure, heterogeneous etc.

LCCC workshop on Cloud Control

22/03/2012 22 / 26

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

## What need to be improved

- From the computer science side:
  - Classification of problems in big classes
  - Standardisation of inputs/outputs/variables for each class
  - Co-design / Control aware software
  - Patience (to explain and to get results)
- From the control theory side:
  - More interest
  - Building a theory that better handles computer science problems
  - Adaptive methods/model free
  - Large scale interconnected systems
- From both side:
  - Spend more time together
  - Spend more time together
- Some inspiring fields
  - Embedded systems (deadline problems, energy optimization, re-allocation, heterogeneous MPSoC, ...)
  - Electrical grids (centralized/decentralized, providers/consumers, cascading failure, heterogeneous etc.)

LCCC workshop on Cloud Control

22/03/2012 23 / 26

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

## What need to be improved

- From the computer science side:
  - Classification of problems in big classes
  - Standardisation of inputs/outputs/variables for each class
  - Co-design / Control aware software
  - Patience (to explain and to get results)
- From the control theory side:
  - More interest
  - Building a theory that better handles computer science problems
  - Adaptive methods/model free
  - Large scale interconnected systems
- From both side:
  - From both side:
  - Spend more time together
  - Some inspiring fields
  - Embedded systems (deadline problems, energy optimization, re-allocation, heterogeneous MPSoC, ...)
  - Electrical grids (centralized/decentralized, providers/consumers, cascading failure, heterogeneous etc.)

LCCC workshop on Cloud Control

22/03/2012 23 / 26

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

## What need to be improved

- From the computer science side:
  - Classification of problems in big classes
  - Standardisation of inputs/outputs/variables for each class
  - Co-design / Control aware software
  - Patience (to explain and to get results)
- From the control theory side:
  - More interest
  - Building a theory that better handles computer science problems
  - Adaptive methods/model free
  - Large scale interconnected systems
- From both side:
  - From both side:
  - Spend more time together
  - Some inspiring fields
  - Embedded systems (deadline problems, energy optimization, re-allocation, heterogeneous MPSoC, ...)
  - Electrical grids (centralized/decentralized, providers/consumers, cascading failure, heterogeneous etc.)

LCCC workshop on Cloud Control

22/03/2012 23 / 26

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

## What need to be improved

- From the computer science side:
  - Classification of problems in big classes
  - Standardisation of inputs/outputs/variables for each class
  - Co-design / Control aware software
  - Patience (to explain and to get results)
- From the control theory side:
  - More interest
  - Building a theory that handles computer science problems
  - Adaptive methods/model free
  - Large scale interconnected systems
- From both side:
  - From both side:
  - Spend more time together
  - Some inspiring fields
  - Embedded systems (deadline problems, energy optimization, re-allocation, heterogeneous MPSoC, ...)
  - Electrical grids (centralized/decentralized, providers/consumers, cascading failure, heterogeneous etc.)

LCCC workshop on Cloud Control

22/03/2012 23 / 26

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

## What need to be improved

- From the computer science side:
  - Classification of problems in big classes
  - Standardisation of inputs/outputs/variables for each class
  - Co-design / Control aware software
  - Patience (to explain and to get results)
- From the control theory side:
  - More interest
  - Building a theory that better handles computer science problems
  - Adaptive methods/model free
  - Large scale interconnected systems
- From both side:
  - From both side:
  - Spend more time together
  - Some inspiring fields
  - Embedded systems (deadline problems, energy optimization, re-allocation, heterogeneous MPSoC, ...)
  - Electrical grids (centralized/decentralized, providers/consumers, cascading failure, heterogeneous etc.)

LCCC workshop on Cloud Control

22/03/2012 23 / 26

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

N. Merchant (eipsse)

gipsLab

LCCC workshop on Cloud Control

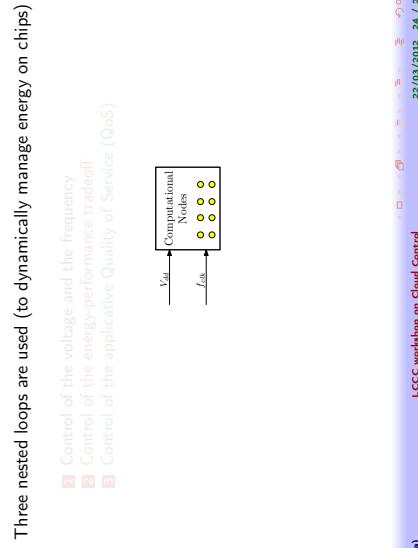
22/03/2012 23 / 26

## gipsa-lab

### What need to be improved

- From the computer science side:
  - Classification of problems in big classes
  - Standardisation of inputs/outputs/variables for each class
  - Co-design / Control aware software
  - Patience (to explain and to get results)
- From the control theory side:
  - More interest
  - Building a theory that better handles computer science problems
  - Adaptive methods/model free
  - Large scale interconnected systems
- From both side:
  - Spend more time together
  - Some inspiring fields
    - Embedded systems (deadline problems, energy optimization, re-allocation, heterogeneous MPSoC, ...)
    - Electrical grids (centralized/decentralized, providers/consumers, cascading failure, heterogeneity, etc.)

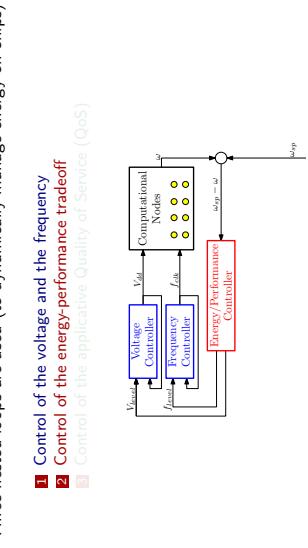
### Feedback loops become essential to handle variability



### gipsa-lab

### Feedback loops become essential to handle variability

- Three nested loops are used (to dynamically manage energy on chips)
- 1 Control of the voltage and the frequency
- 2 Control of the energy-performance tradeoff
- 3 Control of the applicative Quality of Service (QoS)



### gipsa-lab

### Feedback loops become essential to handle variability

- Three nested loops are used (to dynamically manage energy on chips)
- 1 Control of the voltage and the frequency
- 2 Control of the energy-performance tradeoff
- 3 Control of the applicative Quality of Service (QoS)

## Feedback loops become essential to handle variability

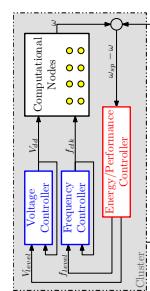


Three nested loops are used (to dynamically manage energy on chips)

LCCC workshop on Cloud Control

N. Merchant

- 1 Control of the voltage and the frequency
- 2 Control of the energy-performance tradeoff
- 3 Control of the applicative Quality of Service (QoS)



Three nested loops are used (to dynamically manage energy on chips)  
Also the approach used in smart grids  
1 Control of the voltage and the frequency  
2 Control of the energy-performance tradeoff  
3 Control of the applicative Quality of Service (QoS)

LCCC workshop on Cloud Control

N. Merchant (eipss)

LCCC workshop on Cloud Control  
22/03/2012 24 / 26

## Feedback loops become essential to handle variability



Three nested loops are used (to dynamically manage energy on chips)  
Also the approach used in smart grids  
1 Control of the voltage and the frequency  
2 Control of the energy-performance tradeoff  
3 Control of the applicative Quality of Service (QoS)

LCCC workshop on Cloud Control

N. Merchant (eipss)

LCCC workshop on Cloud Control  
22/03/2012 24 / 26

## Feedback loops become essential to handle variability

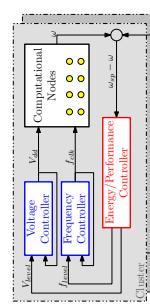


Three nested loops are used (to dynamically manage energy on chips)

LCCC workshop on Cloud Control

N. Merchant

- 1 Control of the voltage and the frequency
- 2 Control of the energy-performance tradeoff
- 3 Control of the applicative Quality of Service (QoS)



Three nested loops are used (to dynamically manage energy on chips)  
Also the approach used in smart grids  
1 Control of the voltage and the frequency  
2 Control of the energy-performance tradeoff  
3 Control of the applicative Quality of Service (QoS)

LCCC workshop on Cloud Control

N. Merchant (eipss)

LCCC workshop on Cloud Control  
22/03/2012 24 / 26

## Feedback loops become essential to handle variability

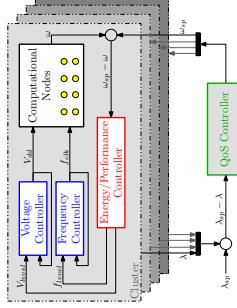


Three nested loops are used (to dynamically manage energy on chips)  
Also the approach used in smart grids  
1 Control of the voltage and the frequency  
2 Control of the energy-performance tradeoff  
3 Control of the applicative Quality of Service (QoS)

LCCC workshop on Cloud Control

N. Merchant (eipss)

LCCC workshop on Cloud Control  
22/03/2012 24 / 26



LCCC workshop on Cloud Control  
22/03/2012 24 / 26

LCCC workshop on Cloud Control  
22/03/2012 24 / 26

## Grenoble Workshop on Autonomic Computing and Control

- Date: 27 may 2014
- Location: Grenoble
- Organisation: Eric Ruttent, INRIA and Stéphane Moraru, Gipsa-lab
- Confirmed speakers:
  - Karl-Erik ARZEN (Lund, Sweden)
  - Alberto LEVA (Milano, Italy)
  - Ada DIACONESCU (Telecom Paris-Tech, France)
  - Suzanne LESFECQ (CEA LETI)
  - Didier DONZEZ (LG)
  - Bogdan ROBU (GIPSA)
  - Eric RUTTEN (INRIA)

## 35th International Summer School of Automatic Control

- Date: September, 8-12, 2014
- Location: Grenoble
- Focus: Modern Tools for Nonlinear Control
- Conclusion



## 35th International Summer School of Automatic Control

- Introduction
- Motivation
- Outline
- Examples
- EB-PID
- Formulation
- Cases
- Simulations
- MapReduce control
- EB-Control
- Conclusion

## 35th International Summer School of Automatic Control

- Introduction
- Motivation
- Outline
- Examples
- EB-PID
- Formulation
- Cases
- Simulations
- MapReduce control
- EB-Control
- Conclusion

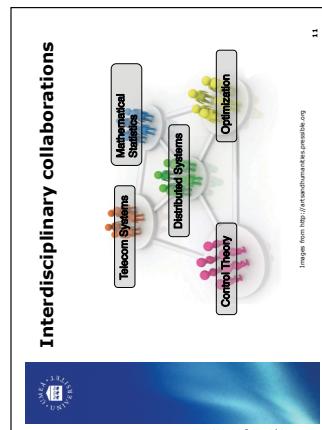
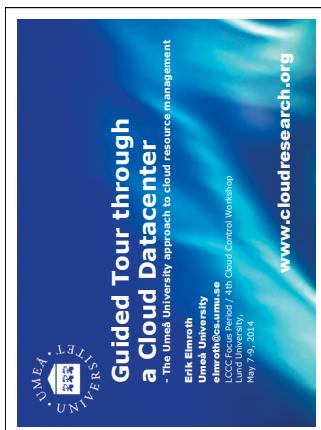
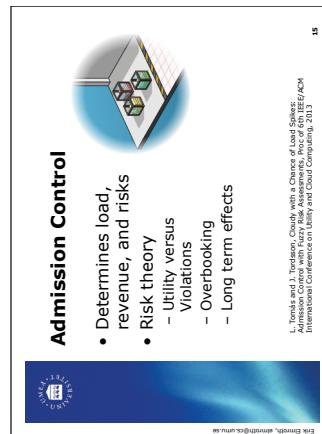
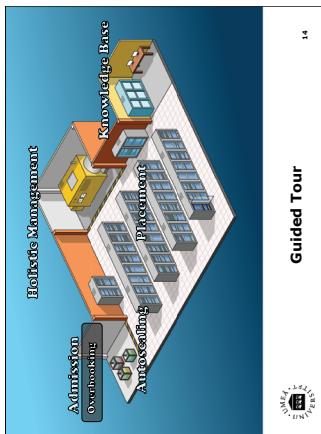


## GUIDED TOUR THROUGH A CLOUD DATACENTER – THE UMEÅ UNIVERSITY APPROACH TO CLOUD RESOURCE MANAGEMENT

Erik Elmroth, Umeå University

By taking a holistic approach to cloud resource management, we aim to transform today's static and energy consuming cloud data centers into self-managed, dynamic, and dependable infrastructures, constantly delivering expected quality of service with acceptable operation costs and carbon footprint for large-scale services with varying capacity demands. The presentation will provide the birds-eye's view of our efforts as well as several glimpses of selected completed, ongoing, and planned research efforts. These efforts address fundamental and inter-twined self-management challenges assuming that there during execution are stochastic variations in capacity need and resource availability, as well as changes in system response and operation costs. Sample challenges include how much capacity to allocate at any time for an elastic application, where to allocate that capacity, if to admit an elastic service with unknown lifetime and future capacity demands, how to optimize the various management tools' concerted actions, etc, while taking into account the need for differentiated quality of service and the scalability requirements of the management tools themselves. For further reading about cloud resource management research at Umeå University, Sweden, please visit [www.cloudresearch.org](http://www.cloudresearch.org).





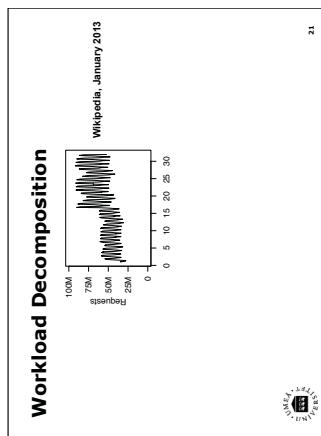
4

## Workload Analysis

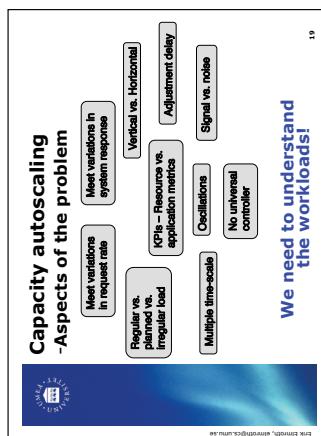
What will your workload look like six years from now?

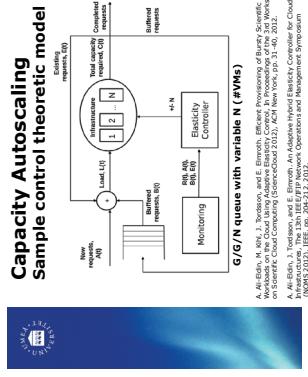
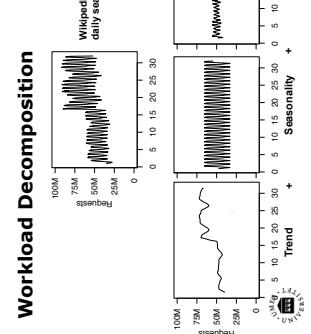
Wikipedia Pages: 1.3 million (2008) 29.6 million (2013)

A. Al-Bilsi, A. Aslani, A. Nefzi, S. Razzouk, S. Stanoić & L. Tuna, O. Simeone, J. Vassilicos and M. Zorzan, "How to Plan for Future Workloads in a Web Environment," *Workload in a Networked World - Proceedings of the 12th IEEE International Conference on Cloud Engineering (IC2E 2014)*, pp. 349-354, 2014.



3

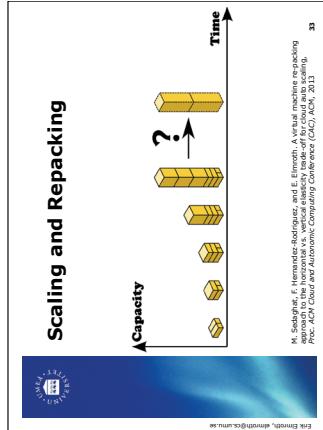
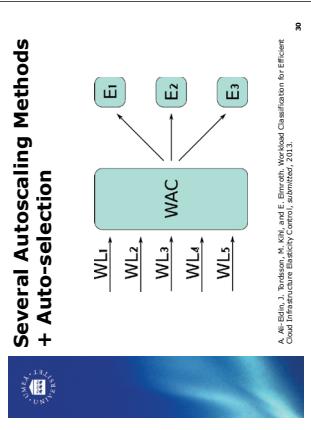




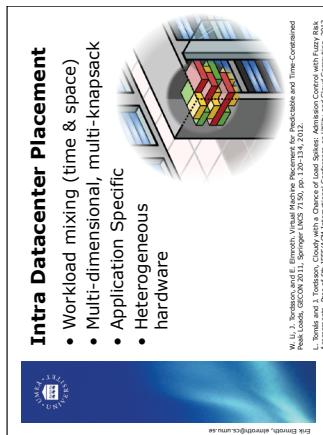
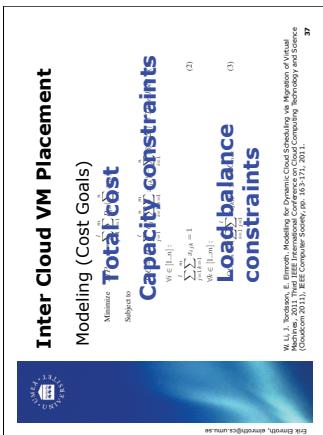
A. Al-Bilsi, M. Kell, I. Zebulum, and E. Elfring, "Efficient Processing of Bursty Scientific Workloads on the Cloud Using Adaptive Resource Control," in Proceedings of the 3rd Workshop on Software-Cloud Convergence (CloudCom 2012), ACM New York, NY, USA, 2012.

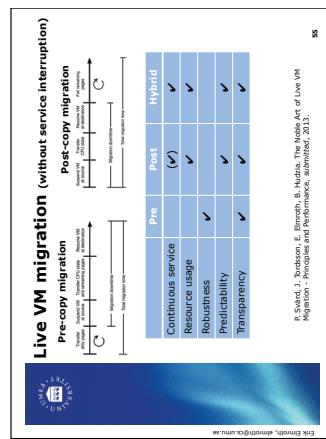
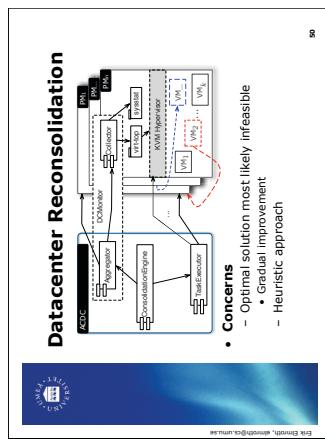
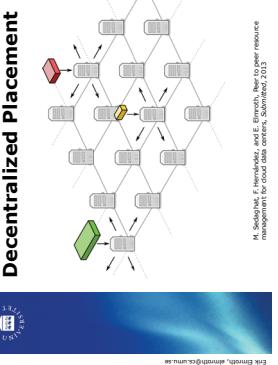
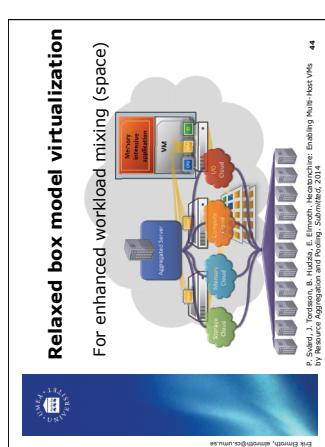
M. Kell, A. Al-Bilsi, I. Zebulum, and E. Elfring, "An Adaptive New Resource Management Strategy for Cloud Workloads," in Proc. 2012 IEEE 21st International Conference on Network Protocols (ICNP 2012), IEEE, pp. 204–212, 2012.

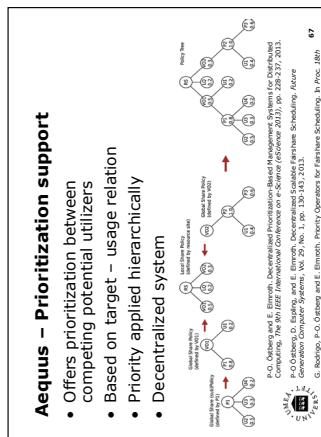
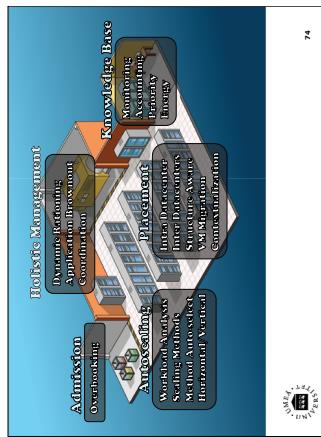
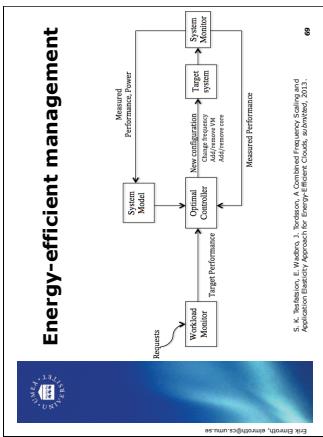
5



6







## Dynamic Resource Rationing

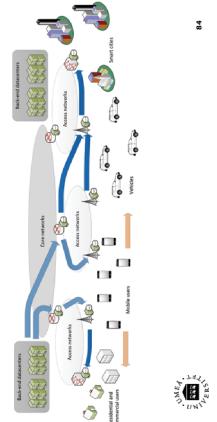
Where to cut when resources are insufficient?

### System Architecture

Two approaches

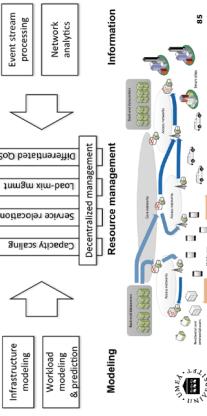
1. Strict QoS-level adherence
2. Overall cost-benefit with QoS-level weights

- Constrained optimization
- Substantial dependency on KPI-type (e.g., latency vs. throughput)
- System feedback on KPI and dimmer effect



## Managing the infinite (or telco) cloud

## Managing the infinite (or telco) cloud

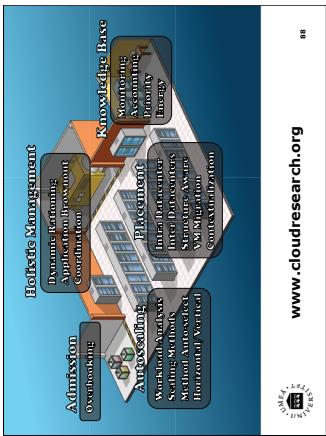


13

[www.cloudresearch.org](http://www.cloudresearch.org)



14









**LUND**  
UNIVERSITY

Department of Automatic Control  
Box 118, 221 00 Lund, Sweden  
[www.lcc.lth.se](http://www.lcc.lth.se)

ISRN LUTFD2/TFRT--7639--SE  
ISSN 0280-5316