Program
9th Cloud Control Workshop

Friiberghs Herrgård, Sweden
June 27 – 29, 2016

Monday June 27th

8:40 Gathering at bus outside Arlanda, Terminal 4

9:00 Bus departure

10:00 **Coffee**

10.30 *Workshop introduction and introductions from participating organizations*
Erik Elmroth, Umeå University, Sweden

11:30 **Keynote: *Container-native cloud platform: devops simplification for diverse workloads***
Malgorzata Steinder, IBM T.J. Watson Research Center, USA

12.20 **Lunch**

Session Chair: Karl-Erik Årzén, Lund University

13.30 *Management of hybrid cloud and application migration*
Azimeh Sefidcon, Ericsson Research, Kista, Sweden

13.50 *Finding the right abstractions for cloud and network virtualization*
Timothy Wood, George Washington University, USA

14:10 *Harness the CPU to increase Containers' security*
Gal Hammer, Red Hat, Israel

14.30

Discussion 1:	Discussion 2:	Discussion 3:	Discussion 4:
Managing decomposed and containerized applications Daniel Espling, Ericsson Research, Sweden	*Network function virtualization: Issues and challenges* Timothy Wood, George Washington University, USA	*Modeling the distributed edge cloud* Karl-Erik Årzén, Lund University, Sweden	*Autoscaling across geo-replicated locations and service tiers* Niklas Carlsson, Linköping University, Sweden

15.30 **Coffee**

Session Chair: Joe Butler, Intel

16.00 *Clouds for crowd management*
Maarten Van Steen, University of Twente, The Netherlands

16:20 *Adaptation strategies for edge clouds*
Omer Rana, Cardiff University, UK

16:40 *Toward hyper-interactive and regulatory-compliant applications in fog computing platforms*
Guillaume Pierre, IRISA, Rennes, France

17:00

Discussion 5:	Discussion 6:	Discussion 7:	Discussion 8:
Getting to a city cloud Maarten Van Steen, University of Twente, The Netherlands	*Distributed cloud control in cyber physical systems* Olov Schelén, Luleå University of Technology, Sweden and P-O Östberg, Umeå University, Sweden	*Improving automation by model-based orchestration* Calin Curescu, Ericsson Research, Sweden	*Software engineering of cloud applications with formal guarantees* Martina Maggio, Lund University, Sweden

18.00 End of session

19.00 **Dinner**

Tuesday June 28th

Session Chair: Maria Kihl, Lund University

8.20 **Keynote: *Feedback in cloud computing systems: Challenges and opportunities in cloud architectures***
Simon Tuffs, Palo Alto Networks, USA

9.10 *Discrete control(s) for autonomic computing*
Eric Rutten, INRIA, Grenoble, France

9.30 *Reconsidering Coherent Distributed Shared Memory for Rack-scale Computing*
Stefanos Kaxiras, Uppsala University, Sweden

9.50 **Coffee**

Session Chair: Olov Schelén, Luleå University of Technology

10.20 *Analytics for network and cloud management*
Rolf Stadler, KTH, Sweden

10:40 *Revising OpenStack internals to operate massively distributed clouds*
Anthony Simonet, INRIA, Nantes, France (replacing Adrien Lèbre, INRIA, Nantes, France)

11:00	Discussion 9:	Discussion 10:	Discussion 11:	Discussion 12:
	Playing a game of Tetris: analytics driven resource and service management Thijs Mesch, Intel, Ireland and Omer Rana, Cardiff University, UK	*Design considerations for the next massively distributed Fog/Edge Cloud Controller* Anthony Simonet, INRIA France	*Rackscale computing: Challenges and Future* Stefanos Kaxiras Uppsala University, Sweden and Ahmed Ali-Eldin, Umeå University, Sweden	*Cloud Robotics: Robots are coming* Giovanni Toffetti, ZHAW, Switzerland and Anders Robertsson, Lund University, Sweden

12.00 **Lunch**

Session Chair: Giovanni Toffetti, ZHAW

13.15 *Adaptive optimal control of cloud Services performance, dependability and costs*
Sophie Cerf, INSA, France

13:30 *Performance comparison of virtual machines and Linux containers*
Zheng Li, Lund University, Sweden

13:45 Short break

14:00 Industry Panel – moderated by David Breitgand, IBM Haifa, Israel

Panelists: Joe Butler, Intel Labs, Ireland
Gal Hammer, Red Hat, Israel
Azimeh Sefidcon, Ericsson Research, Sweden
Malgorzata Steinder, IBM T.J. Watson Research Center, USA
Simon Tuffs, Palo Alto Networks, USA

15:30 **Coffee**

16.00 Social outdoor activities

19.00 **Dinner**

Wednesday June 29th

Session Chair: Geir Horn, University of Oslo

8.30 *Industry, cloud, and 5G- bringing IT forward*
Joakim Persson, Ericsson Research, Lund, Sweden

8:50 *Leveraging memory elasticity in data-parallel task scheduling*
Florin Dinu, EPFL, Switzerland

9.10 *Performance degradation in cloud data centers*
Javid Taheri, Karlstad University, Sweden

9.30 *Benchmarking environment for big data distributed systems: How to make empiric evaluation of (non-) controlled distributed systems easier*
Sara Bouchenak, INSA, France

9.50 **Coffee**

Session Chair: Omer Rana, Cardiff University

10.20 *Increasing the performance of geo-distributed storage systems*
Dejan Kostic, KTH, Sweden

10:40 *Towards energy efficient cloud computing: Facing the big data challenge*
Monica Vitali, Politecnico di Milano, Italy

11:00

Discussion 13:	Discussion 14:	Discussion 15:	Discussion 16:
Resource Scheduling in cloud computing: Issues and research challenges Florin Dinu, EPFL, Switzerland and David Breitgand, IBM Haifa, Israel	*Challenges of Cloud Energy-Aware Applications in an IoT World* Monica Vitali, Politecnico di Milano, Italy	*Resiliency, efficiency and predictability in cloud computing: Existing approaches and open challenges* Sara Bouchenak, INSA, France and Tommaso Cucinotta, Scuola Superiore Sant'Anna, Italy	*To replicate or not to replicate* Cristian Klein, Umeå University, Sweden

12.00 **Lunch**

Session Chair: Ernst Gunnar Gran, Simula Research

13.15 *DieHard: reliable scheduling to survive correlated failures in cloud data centers*
Mina Sedaghat, Umeå University, Sweden

13:30 *Challenges and opportunities with the distributed cloud.*
Victor Millnert, Lund University, Sweden

13.45 *Energy consumption and performance trade-offs on large scale platforms*
Daniel Balouek-Thomert, Ecole Normale Supérieure de Lyon, France

14:00 *Hybrid fault tolerant mechanism for virtual machines*
Abel Souza, Umeå University, Sweden

14:15 *How Google works*
Daniel Pettersson, Google, Sweden

14:30 *Closing, including final discussions and a chance to finalize any discussions (also small groups)*
Erik Elmroth, Umeå University, Sweden

15.00 **Coffee**

15.30 **Bus leaves**

Presentation and discussion topics in order of appearance
7th Cloud Control Workshop
June 9 – 11, 2015

Monday, June 27th

10.30 – 11.30 *Workshop introduction and introductions from participating organizations*
Erik Elmroth, Umeå University, Sweden

11.30 – 12:20 Keynote: *Container-native cloud platform: devops simplification for diverse workloads*
Malgorzata Steinder, IBM T.J. Watson Research Center, Yorktown Heights, NY, USA

Docker is shaping the cloud industry with its container technology. It is helping developers build microservices, move them on the cloud and manage them through the entire DevOps lifecycle. Containers also help cloud providers gain visibility into applications. Transparency of images and containers allows the cloud provider to ensure applications are executing in a secure manner, adhere to compliance rules, and are handled by the cloud management system appropriately for the workload type. A container cloud, where a container is a primary means to deliver applications and virtualize infrastructure, offers an unprecedented opportunity to deliver simplification of management without the loss of flexibility—a game changer, indeed. To capture this opportunity many challenges have to be addressed from platform security and operational efficiency to workload-centric analytics and optimization.

13:30 – 13:50 *Management of hybrid cloud and application migration*
Azimeh Sefidcon, Ericsson Research, Kista, Sweden

The focus of the talk will be on the impact that cloud has had on industry transformation from one side and from the other side, the role that it needs to continue to play in the year to come.

13:50 – 14:10 *Finding the right abstractions for cloud and network virtualization*
Timothy Wood, George Washington University, USA

Virtualization provides clouds an appealing abstraction layer for offering application-agnostic services such as resource management, reliability, and security. However, the isolation provided by virtualization, combined with the semantic gap of an OS and application-agnostic hypervisor, can limit the effectiveness of such services. As a result, breaking the virtual machine abstraction to allow more information to be exchanged or barriers to be crossed more easily can provide significant benefits in some virtualization scenarios.
This talk will explore our work investigating the tradeoffs when puncturing the virtualization layer abstraction. When managing virtual server resources, we found many cases where a "black-box" approach was highly successful. In contrast, our work building a network function virtualization platform shows that breaking VM isolation is critical for extremely latency sensitive I/O applications. Similarly, container technologies have recently seen massive growth in popularity since they provide the flexibility of virtualization, but a lighter weight abstraction layer based on processes instead of operating systems. This talk will cover this tension, exploring the benefits and drawbacks of isolation and abstraction layers for virtual machines and containers.

14:10 – 14:30 *Harness the CPU to increase Containers' security*
Gal Hammer, Red Hat, Israel

The growth in popularity of Containers in the cloud world pushes aside the virtual machines, and with it the benefits gained by using virtual-machine technologies (i.e. AMD-V and Intel-VT). In this talk, a new model will be presented, KVM Containers. The model will be explained as well as some of the challenges ahead.

14:30 – 15:30 Discussion 1: *Managing decomposed (and containerized) applications*
Daniel Espling, Ericsson Research, Sweden

Containers provide a unified and (mostly) technological neutral way to decompose applications into smaller building blocks. These building blocks in turn can thrive in the Microservice-universe and harvest the associated benefits such as independent component life cycles and dedicated software dependencies. However, decomposed applications mean new challenges both for the underlying platform, e.g. in terms of dynamically re-composing the application on top of the set of available services, and for the developers in order to ensure that each permutation of the available services can successfully be assembled into a working instance of the application. In this session; we will discuss the roles of the developer and the platform in terms of managing decomposed services; which responsibilities that are to be handled by each role; and what parts of the process that are possible to automate in order to reduce the operational burden on both the platform and the developers.

14:30 – 15:30 Discussion 2: *Network Function Virtualization(NFV): Issues and challenges*
Timothy Wood, George Washington University, USA

Network Function Virtualization (NFV) is an emerging technology that promises to bring efficient and customizable processing to the network data plane. This has the potential to change networks from "dumb" pipes that simply transport bits, to complex processing pipelines that inspect and transform data as it moves through them. NFV can be used in many areas: in ISP networks for managing VPN gateways, in edge clouds for caching, and in data centers for traffic optimization. This session will give background on NFV and discuss how its use can transform how and where data is processed in both cloud data centers and wide area networks.

14:30 – 15:30 Discussion 3: *Modeling the distributed edge cloud*
Karl-Erik Årzén, Lund University, Sweden

The next generation distributed edge cloud opens up for new types of mission-critical low latency applications, e.g., in robotics and automation, or autonomous vehicles. It is likely that these applications will be on a format that is closer to traditional data-flow models of computations than what is typically currently executed in the cloud. The mission-critical cloud, however, will only be reality if it is possible to provide guarantees on the latency. In order to provide this, one needs suitable models and the associated theoretical machinery. Models are need both for the network, including the nodes, and the communication link, and for the applications. The topic of this discussion session includes what models types that are needed, how to handle the heterogeneity, how to develop local rather than centralized approached, and how to support the required the dynamicity.

14:30 – 15:30 Discussion 4: *Autoscaling across geo-replicated locations and service tiers*
Niklas Carlsson, Linköping University, Sweden

Most cloud elasticity work thus far has been focused on scaling a service deployed in a single location and often by considering a single application tier only. However, some services (e.g., content delivery services delivering globally popular contents) are likely to be replicated across many locations and modern applications are commonly deployed as multi-tier or microservice architectures. In such multi-tiered, multi-location systems a number of interesting questions arise related to how to best perform auto-scaling across service tiers and/or geo-replicated locations. This includes both systems questions related to how to best design decentralized elasticity controllers (global vs. local optimization, consensus, convergence, etc.), but also modeling related questions regarding how to best model the tradeoffs between performance, cost, latency, etc., which themselves are important when designing optimized controllers and protocols for these systems. For example, what aspects of real systems (e.g., variations in network delays, unreliable compute performance from public clouds, etc.) are most important for a model to capture? Are simple queuing models sufficient to capture some of these aspects, or are more advanced models needed to capture geo-replicated scenarios? Which of these aspects are most likely to impact the accuracy of the models? What elasticity-related resource management actions (scheduling, load balancing, etc.) are most affected in such a scenario? These are some of the questions that we hope to discuss in this session in our joint quest to answer how tomorrow's multi-location auto scaling should best be designed.

16:00 – 16:20 *Clouds for crowd management*
Maarten Van Steen, University of Twente, The Netherlands

Monitoring and analyzing crowd behavior is receiving increasingly more attention. In this brief talk I will discuss a recent Dutch initiative to provide an overall framework for gathering data on crowds from WiFi and other types of scanners, and to subsequently analyze the data using a federation of cloud-based solutions. The underlying idea is that there are so many different approaches that working toward a one-size-fits-all solution is not feasible. There are many challenges to face, ranging from attaining accurate high-resolution data on detected devices, to offloading raw or processed data to a cloud, to actual analyses. Meanwhile, an important concern is to maintain privacy. I will briefly present our ideas, but also the problems that we expect to face.

16:20 – 16:40 *Adaptation strategies for edge clouds*
Omer Rana, Cardiff University, UK

Over the last few years we have seen three major trends:
(i) increasing capability in edge devices- such as recent growth and interest in Internet-of-Things and realisation that there is significant heterogeneity in the architecture of IoT devices;
(ii) increasingly ability to externally manage data flow, through a programmable control plane in networks - e.g. through Software Defined Networks and Network Function Virtualization;
(iii) increasing interest to place small-scale clouds closer (particularly for latency sensitive applications) to a user.
Each of these three trends adds additional control parameters that could be made use of to support service orchestration and deployment. The "traditional" data centre oriented provisioning therefore needs to change-to enable our cloud system to "stretch" from a data centre to the edge of the network. The aim of this talk is to identify how we can develop more adaptive cloud architectures and deployment models that could leverage on these trends.

16:40 – 17:00 *Toward hyper-interactive and regulatory-compliant applications in fog computing platforms*
Guillaume Pierre, IRISA, Rennes, France

The mobile network industry is investing heavily in fog computing platforms located at the edge of the networks, in immediate proximity to the end users. Instead of treating the mobile operator's network as a high-latency dumb pipe between the end users and the external service providers, fog computing platforms (sometimes also called mobile edge clouds) aim at deploying cloud functionalities within the mobile phone network, inside or close to the mobile access points. However, traditional clouds are composed of a large number of powerful computing resources located in a handful of data centers located relatively far from the end users, fog computing platforms will be composed of myriads of small nodes located in immediate proximity to the end users, but relatively far from each other. The main software engineering paradigms currently used to develop mobile applications will not apply in such environments. I will discuss potential directions to design developer-friendly middlewares for fog computing environment which still allow the system to optimize performance and exploit the infrastructure to the best of its abilities.

17:00 – 18:00 Discussion 5: *Getting to a city cloud*
Maarten Van Steen, University of Twente, The Netherlands

With the Internet-of-Things gradually coming seriously to life, many municipalities struggle with how to come to a "Smart City." An important issue is that there are many vendors willing to bring in their knowledge, expertise, and products, yet a municipality needs to prevent a lock-in, but also needs to handle concerns on data ownership and privacy. One potential solution is to make use of a "City cloud:" a cloud-based solution owned and exploited by a municipality as a trusted party, through which providers and consumers connect, while the municipality ensures openness, privacy, and so on. In this workshop, we will interactively exploit the feasibility and consequences of such a city cloud.

17:00 – 18:00 Discussion 6: *Distributed cloud control in cyber physical systems*
Olov Schelén, Luleå University of Technology, Sweden

Cyber-Physical Systems (CPS) are typically software defined and distributed over many physical devices that communicate and coordinate (locally and globally) through cloud technologies. Solutions often implement distributed monitoring and adaptive control mechanisms using autonomous system techniques, e.g., Monitor-Analyze-Plan- Execute-Knowledge (MAPE-K) loops, and are as such natural fits for cloud control techniques.
The objective in this group is to identify (parts of) CPS systems that are specifically interesting from the cloud control perspective, and to discuss the state of the art and challenges in this field. Discussion topics include:

- System architectures including information and control flow
- Availability/reliability issues
- Contextual information that can be derived from CPS and implications of use of this in control systems
- Theory and models for defining and validating CPS
- Future projections

To structure the discussions, we will use three case studies from potential application fields as examples / starting points:

- Self-braking cars that use contextual traffic information to regulate behavior (e.g., traffic jam or accident information, moose / reindeers on the road alerts)
- Cloud resource and infrastructure allocation/scaling/scheduling
- IT architectures and infrastructure system that adapt to their environments (e.g., fog / edge computing systems that migrate system components based on temperatures and power requirement predictions)
- Mobile device systems that interact with smart environments (e.g., home environment control and support systems)

17:00 – 18:00 Discussion 7: *Improving automation by model-based orchestration*
Calin Curescu, Ericsson Research, Sweden

Orchestration represents the automated configuration, coordination and management of the cloud resources and services. Most of the times this is implemented as running a set of workflows with little autonomous decision-making.
Future cloud platforms may easily span across several domains, due to either geographical or administrative boundaries, different technological areas (datacenter types, networking domains, OSS/BSS domains) or due to different business entities.
By creating clear models of the resources, the applications, and the lifecycle control loops we may enable orchestration to take autonomous decisions that span across the afore-mentioned boundaries.
In this discussion session, let's explore some of the needs and the challenges of improving orchestration.

17:00 – 18:00 Discussion 8: *Cloud Robotics- Robots are coming*
Giovanni Toffetti, ZHAW, Switzerland and Anders Robertsson, Lund University, Sweden

The connection between the physical world and the virtual world has never been as exciting, accessible, and economically viable as today. Sensors, actors and robots are able to deliver many physical services in several scenarios, including industrial production and home automation, elderly care, assisted living, logistics and cooperative maintenance. In isolation, computing capabilities of robots are however limited by embedded CPUs and small on-board storage units. By connecting robots among each other and to cloud computing, cloud storage, and other Internet technologies centered around the benefits of converged infrastructure and shared services, two main advantages can be exploited. First, computation can be outsourced to cloud services leveraging an on-demand pay-per-use elastic model. Second, robots can access a plethora of services complementing their capabilities (e.g., speech analysis, object recognition, knowledge sharing), enabling new complex functionalities and supporting learning. Cloud robotics is a natural extension to the Internet of Things (IoT). Where IoT devices will gather information about an environment to help make smarter decisions, cloud robotics will be able to use this information and act on it. Although there is clear recognition that Cloud access is required to complement robotics computation and enable functionalities needed for robotic tasks (e.g., self-driving cars), it is still unclear how to best support these scenarios.

Tuesday, June 28th

8:20 – 9:10 Keynote: ***Feedback in cloud computing systems: Challenges and opportunities in cloud architectures***
Simon Tuffs, Palo Alto Networks, Santa Clara, CA, USA

As computing has moved to the cloud, the tradeoff between performance, availability and cost has become increasingly complicated. Distributed cloud based architectures are complex, nonlinear, stochastic and time-varying, and are often inherently unstable under heavy load. Companies that move to the cloud need to maximize performance subject to constraints on availability and cost. This presentation surveys some of the challenges involved in designing and operating micro-service cloud systems, and explores constraints to their stability and efficiency. These constraints will be illustrated with examples drawn from Netflix and Life360.

9:10 – 9:30 *Discrete control(s) for autonomic computing*
Eric Rutten, INRIA, Grenoble, France

Discrete control targets the logical and coordination aspects of autonomic systems. It is defined as the supervisory control of discrete event systems (DES) in the control theory community, based on models such as finite state machines or Petri nets. Alternative forms are proposed in AI (planning) or theoretical computer science (program synthesis). It can be coupled with continuous controllers (e.g. in switching or hybrid systems). Applications are amongst others in synchronisation in multi-thread programming, coordination of multiple autonomic loops, reconfiguration control of FPGA-based architectures.

9:30 – 9:50 *Reconsidering Coherent Distributed Shared Memory for Rack-scale Computing*
Stefanos Kaxiras, Uppsala University, Sweden

A coherent global address space in a distributed system enables shared memory programming in a much larger scale than a single multicore or a single SMP. Without dedicated hardware support at this scale, the solution is a software distributed shared memory (DSM) system. However, traditional approaches to coherence (centralized via “active” home-node directories) and critical-section execution (distributed across nodes and cores) are inherently unfit for such a scenario. Instead, it is crucial to make decisions locally and avoid the long latencies imposed by both network and software message handlers. Likewise, synchronization is fast if it rarely involves communication with distant nodes (or even other sockets). We present such a software DSM system, called Argo, which localizes as many decisions as possible and allows high parallel performance with little overhead on synchronization when compared to prior DSM implementations. This approach enables seamless scaling for rack-scale computing and beyond, entirely at user-level, and allows the investigation of interesting questions about caching, prefetching, data distribution, user-space operation, and system management, in relation to modern datacenter workloads.

10:20 – 10:40 *Analytics for network and cloud management*
Rolf Stadler, KTH, Sweden

Over the last 10 years, analytics methods enabled major advances in various fields of system engineering including self-driving cars and image recognition. We discuss to which extent such methods are applicable to network and cloud technologies.

10:40 – 11:00 *Revising Openstack internals to operate massively distributed clouds*
Anthony Simonet, INRIA, Nantes, France (replacing Adrien Lèbre, INRIA, France)

Academics and industry experts are advocating for going from large centralized Cloud Computing infrastructures to smaller ones massively distributed at the edge of the network. However to favor the adoption of such a model, the development of a system in charge of turning such a complex and diverse network of resources into a global Cloud is critical.

In this talk, we introduce the premises of such a system. As opposed to existing approaches that rely on brokering systems, we will introduce a new model of coordinating a significant number of sites by directly revising and extending the OpenStack software platform with self -* and P2P mechanisms.
The goal of this presentation is to simulate offline discussions on the relevance of such a bottom-up approach in comparison with brokering solutions that can be considered as top-down strategies.

11:00 – 12:00 Discussion 9: *Playing a game of Tetris: analytics driven resource and service management*
Thijs Mesch, Intel, Ireland and Omer Rana, Cardiff University, UK

The game Tetris forms a nice metaphor for resource management: customer workloads can be seen as the titles, were the form and shape express the service characteristics e.g. of a Hadoop workload. The provider wants to place those titles/workloads in their playground/landscape optimally over space for optimized total cost of ownership. So the ultimate question becomes "can we through applying analytics teach the resource manager to play the ultimate game of Tetris"?. To achieve this we need to use analytics to learn a) the size and form of the titles which describes the workload and b) how to best place those titles over space and time. During this workshop we want to discuss the data traces needed to run the analysis, the algorithms needed, and ways to best play the game of Tetris so that both provider and the customer benefit.

11:00 – 12:00 Discussion 10: *Design considerations for the next massively distributed Fog/Edge Cloud Controller*
Anthony Simonet, INRIA, Nantes, France

While several projects have been built using “glue” components in the form of broker or orchestration services, we will discuss in this session whether alternatives strategies can be investigated for designing a system specifically for fog/edge-clouds.

Such a system could leverage P2P techniques to natively interact with low-level mechanisms available on physical resource (compute storage and network) of geographically distributed micro/nano datacenters. In this context, we propose to deal with questions such as the pros and cons of brokering approaches in comparison to a unified system in charge of operating disctinct data centers. How could such a system deal with the velocity of the Cloud Computing eco-system when new technologies appear everyday (VM, containers, unikernel, ...)? Could standards such as OCCI alleviate the complexity of using such software platforms? Finally, we will discuss what extensions could be added to current cloud computing stacks to take the advantage of such massively distributed cloud computing infrastructures.

11:00 – 12:00 Discussion 11: *Rackscale computing: Challenges and Future*
Stefanos Kaxiras, Uppsala University, Sweden and Ahmed Ali-Eldin, Umeå University, Sweden

During the coming few years, we expect to see Rack-Scale Computers (RSC) with thousands of cores, petabytes of solid-state memory, and high-bandwidth / low-latency internal fabrics become a

reality. Such an advancement will allow building custom servers on the fly, where almost the whole infrastructure is software defined. This poses interesting research questions related to server design, the network stack, language runtime systems, and systems software.
In this discussion, We will try to answer questions such as, How should operating systems manage and schedule applications on RCS? Should that OS be something like Xen/KVM, something like Mesos, or something different? How do we control performance, QoS, and overbooking? What is the role of virtualization and are current virtualization solutions sufficient? What problems are best solved in hardware versus software. Are the current directions towards RSC adequate? Can we design RSC to be controllable from the very beginning?

11:00 – 12:00 Discussion 12: *Software engineering of cloud applications with formal guarantees*
Martina Maggio, Lund University, Sweden

The software engineering community has devoted more attention in time to formal guarantees on how software adapts to the current operating conditions. Control theory has been identified as one viable solution for designing the adaptation strategy, together with machine learning, probabilistic model checking and other alternatives. The discussion will be developed along three main lines. First, we will touch upon what type of guarantees should be provided and what kind of guarantees are reasonable to provide. Second, we will discuss how to test and assess these guarantees. Third, we will discuss what the currently used strategies are providing and what are the research gaps that should be filled.

13:15 – 13:30 *Adaptive optimal control of cloud services performance, dependability and costs*
Sophie Cerf, INRIA, France

Many research have been conducted either to improve the dependability or to increase performance of cloud services, for instance using fault-tolerance solutions or through task scheduling techniques.
We investigated an optimization-based solution to control cloud systems in order to provide guarantees in terms of both performance and dependability while reducing utilization costs. We follow a control theoretical approach with special application to the MapReduce framework. Moreover, we aim to be robust to changes in the system and in it's environment by adapting the control algorithm on-line to those changes. Major theoretical challenges arise when combining system adaptation and optimal control to take the best of both approaches.

13:30 – 13:45 *Performance comparison of virtual machines and Linux containers*
Zheng Li, Lund University, Sweden

The current virtualization solution in the Cloud widely relies on hypervisor-based technologies. Along with the recent popularity of Docker, the container-based virtualization starts receiving more attention for being a promising alternative. Since both of the virtualization solutions are not resource-free, their performance overheads would lead to negative impacts on the quality of Cloud services. To help fundamentally understand the performance difference between these two types of virtualization solutions, we use a physical machine as a baseline to investigate the performance overhead of a standalone Docker container against a standalone virtual machine (VM). With findings contrary to the related work, our evaluation results show that the virtualization performance overhead could vary not only on a feature-by-feature basis but also on a job-to-job basis. Although the container-based solution is undoubtedly lightweight, the hypervisor-based technology does not come with higher performance overhead in every case. For example, Docker containers particularly exhibit lower QoS in terms of storage transaction speed.

Wednesday, June 29th

8:30 – 8:50 *Industry, cloud, and 5G - bringing IT forward*
Joakim Persson, Ericsson Research, Lund, Sweden

Much of publicity around the use of cloud technologies currently focus on (personal) services, storage, smart homes, and some more visionary applications like fleets of self driving cars. There is, however, another lesser known but equally interesting field now emerging concerning industrial cloud applications. Clearly, the possibilities that are recognised for the internet-of-things in general, are also applicable to a factory plant where manufacturing tools and machines interact with literary thousands of sensors and actuators. Should all of these become connected, new opportunities arise. With the upcoming 5G technologies, bringing this fabric of data sources and sinks online in a secure and economical way is not an utopia anymore. Moving monitoring and control of the production line to the cloud is desirable for several reasons: convenience, flexibility, predictive maintenance, life-cycle management, security, availability, quality control, and much more. But bringing the monitoring and control to the cloud is difficult for at least as many and tough reasons: timing issues, reliability, security, perceived loss of control, (lack of) trustworthiness, and conservative management adapting to new but unfamiliar technologies at glacier speeds. Despite the difficulties, Ericsson sees an increased interest in this area and there are many research questions to be answered.

8:50 – 9:10 *Leveraging memory elasticity in data-parallel task scheduling*
Florin Dinu, EPFL, Switzerland

In this work we introduce memory elasticity for data parallel tasks and show how to leverage it for improved cluster scheduling. We define memory elasticity as the property of a task to execute correctly and with only a moderate performance penalty, when allocated less memory than it would ideally want. We show that memory elasticity is prevalent among Hadoop mappers and reducers and also exists in the popular Apache Spark, Tez and Flink big data frameworks. Additionally, for Hadoop tasks we build simple models that can accurately predict the performance penalties of less-than ideal memory allocations.
We present ELASTIC, a new cluster scheduler that leverages memory elasticity. ELASTIC may schedule a task even when less memory is available than requested, essentially trading off moderate increases in task runtime for reduced task waiting time and thus providing better overall job completion time. In contrast, current schedulers, such as, e.g., YARN, run tasks only when the amount of memory they requested is available. ELASTIC carefully weighs when to leverage memory elasticity. It achieves this by running lightweight, discrete-event simulations to predict the future state of memory usage in the cluster, and uses the outcome of that simulation to decide whether it operates in greedy mode (schedule tasks immediately even if not enough memory is available) or in regular mode (wait until enough memory is available). By extensive measurements, we show that for many applications both mappers and reducers are elastic in their memory demands, with only moderate penalties for less than optimal memory allocations. Given 10% of ideal memory, Wordcount and PageRank Hadoop reducers are only 1.8x and 1.2x/1.75x slower than ideal. The largest penalty encountered for Hadoop mappers was 1.5x. On a 50 node Hadoop cluster ELASTIC improves on YARN by up to 48% in average job completion time. Extensive simulations show even larger improvements vs YARN. We also show in simulation that ELASTIC can achieve improvements beyond the reach of elasticity-agnostic schedulers.

9:10 – 9:30 *Performance degradation in cloud data centers*
Javid Taheri, Karlstad University, Sweden

Virtualized Data Centers are packed with numerous web and cloud services nowadays. In such large infrastructures, providing reliable service platforms depends heavily on efficient sharing of physical machines (PMs) by virtual machines (VMs). Appropriate consolidation can only be achieved when performance degradation of co-located VMs is correctly understood, modeled, and predicted. To date, because the true negative effect of consolidation is still unknown and unpredictable for most general purpose VMs/services, PMs still remain under-utilized (about 20% of their full capacity) to guarantee the performance of their hosted VMs.
In this presentation, I will introduce novel techniques to measure and quantify performance degradation of VMs on consolidated environments. I will show results of ~1200 hours of profiling and testing of the proposed solutions on two private clouds. Results are promising and prove the efficiency of the proposed solution. At the end, I will introduced open issues and ways to further optimize procedures to gauge performance degradation of in-service VMs.

9:30 – 9:50 *Benchmarking environment for big data distributed systems: How to make empiric evaluation of (non-)controlled distributed systems easier*
Sara Bouchenak, INSA, France

Several cloud computing environments are proposed for developing and executing Big Data applications, distributed data-intensive and compute-intensive applications. MapReduce is a popular programming model for distributed data processing. Extensive research has been conducted on controlling and providing guarantees on the performance and the reliability of MapReduce distributed systems. However, realistic benchmarks are still missing to empirically analyze and compare the effectiveness of these proposals. To date, most MapReduce fault-tolerance and performance improvement solutions have been evaluated using micro-benchmarks in an ad-hoc and overly simplified setting, which may not be representative of real-world applications.
This talk presents MRBS, a comprehensive benchmarking environment for the empirical evaluation of the performance and dependability of MapReduce systems. MRBS includes five benchmarks covering several application domains and a wide range of execution scenarios such as data- intensive vs. compute-intensive applications, or batch applications vs. on-line interactive applications. MRBS allows to inject various types of faults at different rates. It also considers different application workloads and dataloads, and produces extensive reliability, availability and performance statistics. We illustrate the use of MRBS with Hadoop clusters running on Amazon EC2, and on a private cloud.

10:20 – 10:40 *Increasing the performance of geo-distributed storage systems*
Dejan Kostic, KTH, Sweden

Modern distributed systems are geo-distributed for reasons of increased performance, reliability, and survivability. At the heart of many such systems, e.g., the widely used Cassandra and MongoDB data stores, is an algorithm for choosing a closest set of replicas to service a client request. Dynamically changing network conditions pose a significant problem, with suboptimal replica choices resulting in reduced performance due to increase response latency. In this talk I will present our vision for reducing the latency of responses in geo-distributed storage systems. First, I will briefly present GeoPerf, a tool that tries to automate the process of systematically testing the performance of replica selection algorithms for geo-distributed storage systems. Second, I will outline EdgeVar, our system capable of distinguishing between latency changes due to routing changes and congestion build up. Finally, I will describe our newly started Time-Critical Clouds project.

10:40 – 11:00 *Towards energy efficient cloud computing: Facing the big data challenge*
Monica Vitali, Politecnico di Milano, Italy

Energy-awareness and Energy Efficiency are becoming key drivers in Cloud Computing environments. These topics have been widely studied, but classic approaches (e.g. consolidation) have become obsolete with the introduction of more efficient devices at the physical layer. Given these premises, the attention is shifting to the virtual and application layers, in which there is still room for improvement. Resource monitoring, optimization criteria, adaptive mechanisms are being intensively studied from constantly improving efficiency and reducing the environmental impact. In this scenario, the Internet of Things can be used to enrich classical monitored data with external information. Self-adaptation techniques exploit the big amount of information coming from several sources to optimize energy efficiency without neglecting quality of service. However, the management and analysis of this information comes at a cost, and again we have to find a trade-off between the data utility and the cost in terms of energy and performance of exploiting this utility.

11:00 – 12:00 Discussion 13: *Resource scheduling in cloud computing: Issues and challenges*
Florin Dinu, EPFL, Switzerland and David Breitgand, IBM Haifa, Israel

A cloud is a large and complex distributed system, which is characterized by extreme heterogeneity, uncertainty, dynamicity and geographical dispersion of workloads, resources and end users. The workloads share the infrastructure and have widely different objectives. Hence, allocating cloud resources to workloads is a challenging topic, which explains why it is still high on the research agenda since cloud computing has been first introduced circa 2009. As the field of cloud computing matures, and historically disparate communities, such as HPC, BigData, and services converge, what are the next challenges in cloud resource scheduling that we face in the academic and industrial research? The topics we will be discussing include important trade-offs that are shaping innovation in the cloud: good/predictable application performance vs efficient resource utilization or scheduling quality vs latency.

11:00 – 12:00 Discussion 14: *Challenges of Cloud Energy-Aware Applications in an IoT World*
Monica Vitali, Politecnico di Milano, Italy

Traditional Energy-awareness of applications is based on information about data center, servers, and Virtual Machines, collected through a monitoring system. The information coming from Internet of Things (IoT) can complete the information provided by the monitoring system, increasing the awareness about the application behavior, and setting a background information towards application adaptation. This big opportunity comes at a cost: how IoT data should be collected, managed, and analyzed? Which is the cost of managing these data? What can be done to account CO2 emissions on an application basis in a such complex system? This discussion session will focus on challenges raised by IoT and Big Data in the context of Green Applications, trying to understand if traditional approach can be mapped in such an extended domain and how it can be done.

11:00 – 12:00 Discussion 15: *Resiliency, efficiency and predictability in cloud computing: Existing approaches and open challenges*
Sara Bouchenak, INSA, France and Tommaso Cucinotta, Scuola Superiore Sant'Anna, Italy

Building cloud computing services that are able to meet precise resiliency and performance requirements presents a number of challenges that have been addressed at various levels in academia and industry. However, many problems are still unsolved, spanning across the multi-faceted issues involved in the design, realization and operation of such systems. These comprise: meeting a variety of heterogeneous and often contrasting set of requirements, such as availability, fault-tolerance,

security, efficiency, predictable execution, cost; by adopting a variety of techniques including modeling, adaptation and control, with more or less formal underpinning within control theoretical approaches; deploying such techniques at various layers within the cloud provisioning model (IaaS, PaaS, SaaS), implementing them within mechanisms residing at different layers of the software stack, from the OS/hypervisor layer (e.g., scheduling) up to the middle-ware and application layers, and touching on data workload distribution and load balancing.
In the 9th Cloud Control Workshop, we aim for a fruitful discussion among participants centered around the above topics, with particular emphasis on sharing the experience and the individual visions about the challenges lying ahead in this area, as well as focusing on the methodologies for a proper evaluation and effective comparison of the various techniques, where adopting common approaches, benchmarking and simulation environments would be beneficial to ensure the use of a common meter across different research groups.

11:00 – 12:00 Discussion 16: *To replicate or not to replicate*
Cristian Klein, Umeå University, Sweden

This discussion session is about replicating experimental results in computing systems research. Commonly mentioned reasons why replication is desirable are: (a) ability to independently confirm a research finding by another group; (b) ability to test the applicability of a finding in a broader context, e.g., different parameters, hardware or software; (c) ability to do an apples-to-apples comparison of a new system with an already studied one; (d) reusing experimental artifacts for a new research question. On the other hand, some people think that replication is pointless, as it risks repeating the mistakes in the original experiment and falsely confirm a wrong results.
The session starts with a short presentation on different levels for obtaining the same result repetition, replication, variation, reproduction and corroboration and the growing importance that repeatability receives in various research communities. Then participants are asked to share their experience and discuss along the following points:
* How important is repeatability in computing systems research?
* What are the current (technical or human) barriers to repeatability?
* What positive and negative experiences did you have with repeating somebody else's results?
* How would you like to receive an experiment?
* Do you aim for your results to be repeatable? What did you do in the past to enable this?

13:15 – 13:30 *DieHard: reliable scheduling to survive correlated failures in cloud data centers*
Mina Sedaghat, Umeå University, Sweden

In large scale data centers, a single fault can lead to correlated failures of several physical machines and the tasks running on them, simultaneously. Such correlated failures can severely damage the reliability of a service or a job.
In this talk, we discuss the the impact of stochastic and correlated failures on job reliability in a data center. We present a statistical reliability model and an approximation technique for computing a job's reliability in the presence of correlated failures. In addition, we address the problem of scheduling a job with reliability constraints. We formulate the scheduling problem as an optimization problem, with the aim being to achieve the desired reliability with the minimum number of extra tasks. Finally, we present a scheduling algorithm that approximates the minimum number of required tasks and a placement to achieve a desired job reliability.

13:30 – 13:45 *Challenges and opportunities with the distributed cloud.*
Victor Millnert, Lund University, Sweden

Today we are seeing many new, enabling technologies for controlling things in the Cloud. One of these is containers, enabling us to start/stop new micro-services on a different time-scale than what

was possible with virtual machines. We are also seeing a new open source environment, Calvin, for the development of Internet-of-Things applications. This, together with the distributed infrastructure of the Cloud, brings with it new ways of designing control-strategies for large systems. In this talk I will discuss the possibilities and demands for designing control-systems running in the Cloud, and how these could be translated into demands on the Cloud. An example is tactile feedback, requiring an end-to-end delay of about 1ms, imagine the applications we could have. A different one is the traffic coordination of a country full of autonomous cars.

13:45 – 14:00 *Energy consumption and performance trade-offs on large scale platforms*
Daniel Balouek-Thomert, Ecole Normale Supérieure de Lyon, France

As the demand for Cloud infrastructures increases dramatically across the globe, the reduction of energy consumption in such infrastructures has become a challenge, urging for solutions that concurrently mitigate the environmental impact while maximizing the economic benefits. In this context, we intend to perform an efficient selection of resources with respect to a certain quality of service (mainly throughput of computing tasks) and energy consumption. The implementation of the task-to-resource allocation policy consists in picking in real time at Cloud provider's end the best combination of resources, in order to fit the customer's needs at lowest cost and risk with respect to his preferences, and lowest energy consumption.

This research work synergize two state-of-the-art technologies by combining multi-objective evolutionary algorithms with trade-off mechanisms using the DIET middleware (http://graal.ens-lyon.fr/DIET/), involving in vivo validations over the Grid'5000 testbed.

The New-Generation SR company proposes a multi-layer solution named NUVEA, based on the industrial transfer of this work. NUVEA (http://nuvea.eu/) is divided into 3 services: AUDIT, OPTIMIZATION and BROKER and aiming at optimizing the yield management of data centres while giving more flexibility to cloud consumers. The existing NUVEA platform allows cloud customers to permanently measure, take over, manage, benchmark and optimize their distributed IT infrastructure at the web- and enterprise-level in a highly flexible and real-time manner.

14:00 – 14:15 *Hybrid fault tolerant mechanism for virtual machines*
Abel Souza, Umeå University, Sweden

High availability and fault tolerance are crucial aspects of today's services, as mission critical applications need to keep running even in case of failures. Currently, most hypervisors' fault-tolerant systems use application agnostic solutions based on software replication techniques, where a standby virtual machine closely tracks the execution of the primary and takes over upon failure.
These techniques are often implemented by periodic state checkpointing (active/passive replication). They may, yet, suffer from excessive overhead under different circumstances. The novel COarse Grain LOck Stepping (COLO) approach addresses this problem by reducing the number of checkpoints by triggering on-demand VM replications only when the two replicas differ from the user point of view (active/active replication). However, in some cases COLO has shown to increase the frequency of checkpoints, leading to performance degradations and to an inefficient use of resources. Hence, there is no solution that presents a good performance in all the cases. In this talk, we will briefly present a control theoretic approach that dynamically switches between the COLO and checkpointing techniques in order to minimize the number of checkpoints

14:15 – 14:30 *How Google works*
Daniel Pettersson, Google, Sweden

(On short notice. No abstract.)