Program
11th Cloud Control Workshop

Haga Slott, Enköping, Sweden
June 12 – 14, 2017

Monday June 12th

8:30 Gathering at bus outside Arlanda, Terminal 4

8:40 Bus departure

9:50 **Coffee**

10.20 *Workshop Introduction and Participant Presentations*
Erik Elmroth, Umeå University, Sweden

11:10 **Keynote: *Mechanism Design and Implementation for Efficient and Trustworthy Cloud Markets***
Azer Bestavros, Boston University, MA, USA

12.00 **Lunch**

Session Chair: Karl-Erik Årzén, Lund University, Sweden

13.10 *Function-as-a-Service: Use Cases, Promises, Challenges, and Research Directions*
David Breitgand, IBM Haifa Research Lab, Israel

13.30 *Resource and Workload Modelling in Highly Distributed Infrastructures*
Martin Körling, Ericsson, Stockholm, Sweden

13:50 *Scaling Resource Allocations for Data-Processing Frameworks*
Dick Epema, TU Delft, Netherlands

14.10 *Using Several Clouds to Improve Storage Dependability: From DepSky to SCFS and Beyond*
Miguel Correia, INESC-ID/IST, Portugal

14.30

Discussion 1 (Manegen):	Discussion 2 (Vagnslidret):	Discussion 3 (Falkenb.)	Discussion 4 (Natt & Dag):
Development and Deployment of Industrial IoT Applications using Edge Cloud and Deep Network Integration Johan Eker, Ericsson Research, Lund, Sweden	*What Can the Blockchain Solve that Distributed Databases Cannot?* Dick Epema, TU Delft, Netherlands and Claus Pahl, Free University of Bozen-Bolzano, Italy	*Performance Engineering for Microservices: Research Challenges and Directions* André van Hoorn, Universität Stuttgart, Germany	*Dependable Modelling and Verification for Increasingly Dynamic Cyber Physical Systems* John. M Kennedy and Thijs Metsch, Intel labs, Leixlip, Ireland

15.40 **Coffee**

Session Chair: Maria Prandini, Politecnico di Milano, Italy

16.10 *Enabling 5G Networks through Red Hat's Open Source Portfolio*
Luis Tomas, Red Hat, Spain

16:30 *Decoupled Access-Execute: Pioneering Compilation for Energy-Efficiency*
Alexandra Jimborean, UPMARC, Uppsala University, Sweden

16:50

Discussion 5 (Manegen):	Discussion 6 (Vagnslidret):	Discussion 7 (Falkenb.)	Discussion 8: (Natt & Dag)
Data-Driven Cloud Engineering and Management Daniel Gillblad, Swedish Institute of Computer Science and Rolf Stadler, KTH, Sweden	*Function-as-a-Service (FaaS): Research Challenges and Open Issues* David Breitgand, IBM Haifa Research Lab, Israel	*Understanding the Characteristics of Cloud Applications* Ahmed Ali-Eldin, Umeå University, Sweden and Alexandra Jimborean, UPMARC, Uppsala University, Sweden	*Modelling Application Data for Analysis* Geir Horn, University of Oslo, Norway

18.00 End of session

19.00 **Dinner**

Tuesday June 13th

Session Chair: Uwe Schwiegelshohn, TU Dortmund, Germany

8.20 *IAM on Hybrid Cloud: Next Generation Security Model to Create an Interoperable Cloud*
Jeyakeerthi Jeyappragash, Twitter, US

8.50 *Implementing and Evolving a Rack Scale Design based SW Defined HW Infrastructure*
Tomas Fredberg, Ericsson, Stockholm, Sweden

9.10 *The Power of (Partial) Redundancy for Curbing Tail Latency*
Lydia Chen, IBM Zürich, Switzerland

9.30 *Autonomous Strategies for Elastic Data Stream Processing in the Fog: from Centralized to Distributed Approaches*
Francesco Lo Presti, University of Rome, Italy

9.50 **Coffee**

Session Chair: Gal Hammer, Red Hat, Israel

10.20 *AI: Beyond Machine Learning for Management*
Daniel Gillblad, Swedish Institute of Computer Science, Sweden

10:40	Discussion 9 (Manegen):	Discussion 10 (Vagnslidret):	Discussion 11 (Falkenb.):	Discussion 12 (Natt & Dag):
	Efficient Workload Modeling for future Software Defined Infrastructures Tomas Fredberg, Ericsson, Stockholm, Sweden	*Managing the Trade-Off between Power and QoS of Applications: an Old Story in a New Heterogeneous World of Things* Jakub Krzywda, Umeå University, Sweden and Monica Vitali, Politecnico di Milano, Italy.	*Making the Cloud Extremely Dependable* Miguel Correia, INESC-ID/IST, Portugal	*The Grand Cluster Resource Utilization Challenge* Marcel Blöcher, Technical University of Darmstadt, Germany

12.00 **Lunch**

Session Chair: Frédéric Desprez, Inria, Grenoble, France

13:00 *Enos: a Holistic Framework for Conducting Scientific Evaluations of OpenStack*
Adrien Lebre, Inria / IMT-Atlantique, Nantes, France

13:20 *Checkpointing In-Memory Data Analytics Applications with Panda*
Bogdan Ghit, TU Delft, Netherlands

13:35 *Model-free QoS Control Using Cooperative Fuzzy Reinforcement Learning*
Olumuyiwa Ibidunmoye, Umeå University, Sweden

13.50 *Short break*

14:00 Panel Discussion – moderated by Johan Eker, Ericsson Research, Lund, Sweden

Panelists:

- Azer Bestavros, Boston University, MA, USA
- Joe Butler, Intel Labs, Ireland
- Lydia Chen, IBM Zürich, Switzerland
- Jeyakeerthi Jeyappragash, Twitter, CA, USA
- Azimeh Sefidcon, Ericsson Research, Stockholm, Sweden
- Steve Webster, Google, CA, USA

15:30 **Coffee**

16.00 Social outdoor activities

19.00 **Dinner**

Wednesday June 14th

Session Chair: Joakim Persson, Ericsson Research, Lund, Sweden

8.20 *ML for Datacenter Cooling Control - Progress and Challenges at Google*
Steve Webster, Google, US

8:50 *Prediction Models for Service Responsiveness of Deep Neural Networks in a Cloud Setting*
Evgenia Smirni, College of William and Mary, US

9.10 *HOPS a Platform for Data Analytics in Data Centers*
Seif Haridi, Swedish Institute of Computer Science and KTH, Sweden

9.30 *Benchmarking Public Clouds With Cloud Workbench - Do You Really Know What You Are Paying For?*
Philipp Leitner, University of Zürich, Switzerland

9.50 **Coffee**

Session Chair: Evgenia Smirni, College of William and Mary, US

10.20 *Taking Elasticity to the Edge*
Srikumar Venugopal, IBM Ireland

10:40

Discussion 13 (Manegen):	Discussion 14 (Gustaf):	Discussion 15 (Falkenberg)	Discussion 16 (Natt & Dag):
Causal Analytics for Automating Data Center Management James Kempf, Ericsson Research, CA, USA	*Data Driven and Distributed Optimization in Cloud Management: Opportunities and Challenges* Alessandro Papadopoulos, Mälardalen University, Sweden and Maria Prandini, Politecnico di Milano, Italy	*Command & Conquer: Control and Assurance in Distributed Systems* Joe Butler and Thijs Metsch, Intel labs, Leixlip, Ireland	*Clouds & 5G: Interplay, Opportunities and Challenges* Dimosthenis Kyriazis, University of Piraeus, Greece and Luis Tomas, Red Hat, Spain

12.00 **Lunch**

Session Chair: Maria Kihl, Lund University, Sweden

13.00 *The Evaporating Cloud*
Guillaume Pierre, University Rennes 1, France

13:20 *Event-Driven Bandwidth Allocation to a Network of Cameras with Formal Guarantees*
Martina Maggio, Lund University, Sweden

13.40 *An Efficient Communication Aware Heuristic for Multiple Cloud Application Placement*
Pedro Silva, Inria, ENS de Lyon, France

13:55 *Dynamic Control of NFV Forwarding Graphs with end-to-end Constraints*
Victor Millnert, Lund University, Sweden

14:10 *Auto-scaling Workflows: A Comparative Study*
Ahmed Ali-Eldin, Umeå University, Sweden

14:25 *Closing*
Erik Elmroth, Umeå University, Sweden

14.40 **Coffee**

15.10 **Bus departure**

16.30 **Latest arrival at Arlanda airport, terminal 4**

Presentations and Discussion Sessions in Order of Appearance
11th Cloud Control Workshop
June 12 – 14, 2017

Monday, June 12th

10.20 – 11.10 *Workshop Introduction and Participant Presentations*
Erik Elmroth, Umeå University, Sweden

11.10 – 12:00 Keynote: ***Mechanism Design and Implementation for Efficient and Trustworthy Cloud Markets***
Azer Bestavros, Boston University, MA, USA

Despite the increased adoption of cloud computing, we do not possess (yet) a good handle on how to define and verify "cloud trustworthiness" -- how to design mechanisms that allow control of attributes related to performance, reliability, security, privacy, and economic utility, and how to expose tradeoffs along these dimensions to cloud customers in ways that are both practical and usable. In this talk, I will summarize research pursued for almost a decade at Boston University, aiming to address these issues. I will start by presenting models for the expressive specification of elastic cloud supply and demand that enable management and control of an efficient, trustworthy cloud marketplace, which I will exemplify by presenting three game-theoretic mechanisms for cloud resource management. The first mechanism enables selfish parties to collocate their workloads in an attempt to minimize the individual costs they incur to secure the shared cloud resources necessary to support their application requirements. The second mechanism enables rational parties to coordinate their use of a shared resource so as to maximize their individual utilities by creating a marketplace for trading usage rights. The third mechanism enables dynamic pricing of IaaS clouds in a way that ensures efficient utilization of providers' resources, while guaranteeing rational fairness to all customers. I will conclude with an overview the Massachusetts Open Cloud (MOC), the first prototypical implementation of an Open Cloud eXchange (OCX) marketplace that makes the adoption of these (and other) trustworthy mechanisms possible in an operational at-scale setting. Time permitting, I will present examples of platforms that we have developed on the MOC, targeting applications ranging from smart-city data mechanics to secure multi-party analytics over private data assets.

13.10 - 13.30 *Function-as-a-Service: Use Cases, Promises, Challenges, and Research Directions*
David Breitgand, IBM Haifa Research Lab, Israel

Recently, Function-as-a-Service (FaaS) has emerged as an attractive new form of PaaS. FaaS are natively supported by most leading cloud providers. Also, several smaller companies offer their cross cloud FaaS platforms on top of traditional IaaS. Initially viewed as a “glue” for service mashup, FaaS is now gradually evolving towards a born in a cloud application server. FaaS is particularly well suited for dealing with IoT scenarios. In this talk, I will outline challenges and research directions related to FaaS from a provider perspective. In this context, I will discuss a problem of joint optimization of heterogeneous cloud resources and long term contractual commitments. Also, I will discuss using FaaS in the edge in the context of a new EU project 5G-MEDIA, dealing with immersive media scenarios in edge computing.

13.30 - 13.50 *Resource and Workload Modelling in Highly Distributed Infrastructures*
Martin Körling, Ericsson, Sweden

Building a distributed cloud infrastructure to support high-data-volume and low-latency applications, for example in 5G, industrial, big data, and IoT scenarios, requires capacity planning as well as dynamic optimization and automation of workload life-cycle management. The principles for a quantitative model for both resources and workloads are outlined. The resource model covers computing, storage as well as networking and the workload model supports elastic operating points. Machine learning and machine decision models can play an important role. Early examples and illustrations are presented and several challenges for further study are identified.

13.50 - 14.10 *Scaling Resource Allocations for Data-Processing Frameworks*
Dick Epema, TU Delft, Netherlands

Single clusters, datacenters, and clouds may deploy multiple instance of multiple data-processing frameworks such as Hadoop, Spark, and many others. These instances may live for a long time, during which workloads with some arrival process are submitted to them. Because the resource demands of these workloads may vary considerably over time, static resource allocations are not appropriate. So the question is how to dynamically adapt the resource allocations to these frameworks. We will address this question both when there is no resource contention (e.g., in clouds with "infinite" resources) when the issue is to right-size the allocations, and when there is contention (e.g., in heavily loaded clusters) when the issue is to balance the resources in a more or less fair way.

This is joint work with Bogdan Ghit who recently obtained his PhD degree at TU Delft, and Aleksandra Kuzmanovska (TU Eindhoven).

14.10 - 14.30 *Using Several Clouds to Improve Storage Dependability: From DepSky to SCFS and Beyond*
Miguel Correia, INESC-ID/IST, Portugal

File storage is one of the most successful use cases for cloud computing. Services like Dropbox, Google Drive, Amazon S3, and Apple iCloud Drive are widely used worldwide to store both personal and professional data. However, there are reasonable concerns about the security and dependability of data stored in such services, with frequent news about outages, espionage, and file corruption. Using several clouds to store data instead of only one is an interesting solution for these problems. The talk will start by introducing DepSky, a cloud storage service that leverages the use of several clouds forming a cloud-of-cloud to provide dependability and security despite accidental and intentional arbitrary ("Byzantine") faults. Then it will present the differences to SCFS, a dependable POSIX-compliant file system based on DepSky. Both systems mask faults using replication. Experimental results with real commercial clouds and current work on the topic will be briefly presented.

14.30 - 15.40 Discussion 1: *Development and Deployment of Industrial IoT Applications Using Edge Cloud and Deep Network Integration*
Johan Eker, Ericsson Research, Lund, Sweden

Highly heterogeneous cloud environments that span from large data centers to edge nodes and even mobile devices offer exciting opportunities in order to address problems around latency, privacy, local redundancy, etc. However, such cloud setups come with additional complexity. The client-

server model has served us well in developing web services, but in order to leverage the capabilities of edge cloud, new programming models are called for. This session will take a developer perspective on edge cloud applications.

14.30 - 15.40 Discussion 2: *What Can the Blockchain Solve that Distributed Databases Cannot?*
Dick Epema, TU Delft, Netherlands and Claus Pahl, Free University of Bozen-Bolzano, Italy

Blockchain is a technology for digitizing trust among multiple parties without requiring central components. It promises a mechanism of smart contracts that invoke programmed actions when certain conditions are satisfied. The question is whether the blockchain is a technology in search of problems, or a solution to problems that we cannot solve otherwise.
This session will investigate this question, in particular for IoT/cloud/edge computing in which the blockchain may address problems like data provenance and mobile software authentication. How far does the smart contract concept stretch, and what are the performance limitations if we have to rely on lightweight devices?

14.30 - 15.40 Discussion 3: *Performance Engineering for Microservices: Research Challenges and Directions*
André van Hoorn, Universität Stuttgart, Germany

Microservices are an emerging software architectural style enabling the efficient use of cloud computing technologies and agile software engineering paradigms and processes, such as DevOps and Continuous Delivery (CD). Performance, particularly including measures such as elasticity and resilience, is promised as a benefit of microservice architectures. However, existing software performance engineering methods, techniques, and tools currently lack explicit support for the characteristics of microservice architectures and related DevOps/CD processes. For example, long-running performance tests conflict with rapid deployments of changes into production as part of CD; performance modeling/prediction has so far focused on upfront capacity planning, which may be less relevant for microservices. The objective of this session is to identify promising research challenges and directions with respect to this topic. A leading questions for this discussion could be “What are performance challenges for microservices in industrial practice?”, “What performance engineering approaches would be helpful to tackle these challenges?”, and “How far are we from providing/using such approaches?” The discussion could start with a short summary of previous discussions/presentations with a slightly different focus, resulting from Dagstuhl seminar 16394 (see also https://doi.org/10.1145/3053600.3053653).

14.30 - 15.40 Discussion 4: *Dependable Modelling and Verification for Increasingly Dynamic Cyber Physical Systems*
John. M Kennedy and Thijs Metsch, Intel labs, Leixlip, Ireland

Multi-Agent Systems and Cyber Physical Systems (CPS) have much in common, however CPS have the added complexity of interactions with their (potentially very dynamic) physical environment. Individual components of the CPS may themselves be highly dynamic and adaptable in nature. Both modelling and the verification of formal and real models in this world become significantly more complex – and significantly more important given the potential for physical harm.

In this session we want to discuss approaches to address these challenges. What are the limitations in emerging digital twin approaches? How can these limitations be tackled? Are self-organizing autonomous particle/swarm based systems more feasible? …

16.10 - 16.30 *Enabling 5G Networks through Red Hat's Open Source Portfolio*
Luis Tomas, Red Hat, Spain

The Superfluidity EU project focuses on 5G networks, and goes one step further into the virtualization and orchestration of different network elements, including radio and network processing components, such as BBUs, EPCs, load balancers, SDN controllers, and others. To take advantage of the mobile/edge network possibilities, as well as to better tackle its associated challenges, we foresee a mix world of VMs and containers, as well as different deployment models (side by side VMs and containers, and nested containers running on top of VMs). To enable this flexibility and allow the instantiation of these services on the fly anywhere on the network (core, aggregation or edge) different components are needed at the different layers defined by ETSI ISG, i.e., VIM, VNFM and NFVO. Superfluidity proposes to use both OpenStack and OpenShift, with Kuryr as the glue enabling both deployment models. At the top layer (NFVO) we propose the use of ManageIQ with extended Ansible support to allow easy on-board, creation and orchestration of new services across different deployments.

16.30 - 16.50 *Decoupled Access-Execute: Pioneering Compilation for Energy-efficiency*
Alexandra Jimborean, UPMARC, Uppsala University, Sweden

Energy efficiency is essential for performance, to prevent thermal hazards, and to enable more simultaneously-active cores. This talk presents software decoupled access-execute (DAE), a compilation technique we pioneered for improving energy efficiency. The compiler decouples the code into coarse-grain memory-bound and compute-bound phases, to re-enable hardware capabilities for energy management, to control data communication vs. data processing and to enhance memory- and instruction-level-parallelism. The end result is 25% energy savings, on average, for memory-bound applications, with negligible impact on performance.

16.50 - 18.00 Discussion 5: *Data-driven Cloud Engineering and Management*
Daniel Gillblad, Swedish Institute of Computer Science and Rolf Stadler, KTH, Sweden

Recent advances in computing, storage, and networking technology enable data analytics on large sets of operational data, such as network and cloud statistics, events and logs. It is often feasible to perform prediction and model computation in real-time and in a distributed fashion.
The discussion group will analyze the impact of these developments on cloud engineering and management and explore questions and topic areas like the following: For which tasks are analytics-based functions superior to “traditional” solutions? What are the success stories? What specifically are the research challenges? What is the potential of deep learning in this context? Should data scientists address these issues, or should they be solved by network and cloud engineers who are trained in data science? Beyond using data science, what is the potential of full AI for future cloud engineering and management?

16.50 - 18.00 Discussion 6: *Function-as-a-Service (FaaS): Research Challenges and Open Issues*
David Breitgand, IBM Haifa Research Lab, Israel

FaaS is the latest major addition to the already complex and diverse cloud computing landscape. Initially introduced by Amazon in 2014, it now rapidly evolves as a popular PaaS flavor with all major cloud vendors offering their versions of FaaS and a few open source projects rapidly developing. FaaS offers several advantages. Among them are fine granularity of billing, seamless elasticity and scalability, high level of abstraction when developing a solution, improved

manageability. FaaS is very natural in event-driven IoT scenarios ranging from smart home to smart factory and smart city. It is very suitable for fast PoC allowing rapid programmable service mashup.

Yet, FaaS is still in its early stage and several important gaps exist. Until recently, FaaS did not deserve sufficient attention from the cloud research community. We believe that there exist a range of timely research opportunities in this field. In this discussion session, we will quickly overview SOTA in FaaS and then present and discuss several research challenges related to this approach. We will be examining both cloud and edge parts of the ecosystem and both cloud provider and cloud consumer perspectives.

16.50 - 18.00 Discussion 7: *Understanding the Characteristics of Cloud Applications*
Ahmed Ali-Eldin, Umeå University, Sweden and Alexandra Jimborean, UPMARC, Uppsala University, Sweden

A recent study on the usage of IoT devices shows that alternating high-intensity workloads (phases) with lower-intensity workloads on the same device, allows the battery to recover and increases battery lifetime. In contrast, the race-to-sleep approach, if used for a very long time, drains the battery entirely. One of the research challenges in IoT is to decide where to process the collected data in order to balance communication and performance bottlenecks. Yet, there are many other factors to account for such as the computing power of devices, costs of communicating data, speed of transmitting data, reliability, etc.

Regarding cloud computing, what are the main challenges, performance bottlenecks, energy demanding applications, vulnerabilities? Can they be addressed by transforming the code?

16.50 - 18.00 Discussion 8: *Modelling Application Data for Analysis*
Geir Horn, University of Oslo, Norway

There are several Domain Specific Languages like TOSCA, CAMEL, and CloudML, for describing application structure that allows efficient and optimised application deployment (mapping). However, how can we model and describe the data sets used and processed by an application? The current approaches largely consider the applications as data parallel where the different application components can be scheduled and executed in independently and in parallel. In reality, the different modules of an application will be operating on the same data sets. This implicitly imposes a work-flow of the application modules. Will the way we model data decide how we eventually will model the application deployment? Will it be sufficient just to model the data processing work flow? Why? How to do the modelling? If the data set is large, and if several applications are processing the same data, can we still model the applications independently? Moving data between data centres can be too costly, so the modules from two applications processing the same data set should perhaps be mapped to the same data centre? In this way the data model will impact both execution modelling, mapping and scheduling. This is a vast and unexplored territory, which cannot be solved in a short discussion session, but the session can bring up ideas for future research to be further explored after the CCW.

Tuesday, June 13th

8.20 - 8.50 *IAM on Hybrid Cloud: Next Generation Security Model to Create an Interoperable Cloud*
Jeyakeerthi Jeyappragash, Twitter, US

The next generation systems are dealing with complexity in multiple dimensions. (i) Systems have to span across in-house servers to datacenters to multiple public clouds; (ii) Deal with various new technology and infrastructures like different types of storage services, different types of querying and analytics solutions and different processing libraries from stream computes to ML. (iii) There is also polyglot of application architectures that are springing up from traditional three-tier apps to fine-grained services.

Given this complexity - capacity management, SLA management and the developer productivity becomes a serious challenge. Each one is a topic of itself, but what we will focus on is a small part of security which is critical to bind this all together.

In order for us to manage and trust the system, we need to solve the problem of security in layers: Service Identity; Credential Management; Authentication; Resource Cataloguing; Contract Management; Authorization; and Enforcement.

There are several components like containers, RPC frameworks, proxies, that allow us to plug-in and provide a secure end-end system. In this talk, we will walk through how each part of this problem is solved by each of these components to create a full on secure system that can be trusted and is reliable.

We will also walk through a reference architecture to accomplish policy and access management at scale.

8.50 - 9.10 *Implementing and Evolving a Rack Scale Design based SW Defined HW Infrastructure*
Tomas Fredberg, Ericsson, Sweden

The story and some learnings about the Ericsson Hyperscale Datacenter System implementation of a Software Defined HW Infrastructure based on the Intel Rack Scale Design architecture.

9.10 - 9.30 *The Power of (Partial) Redundancy for Curbing Tail Latency*
Lydia Chen, IBM Zürich, Switzerland

Offering consistent low latency remains a key challenge for distributed applications, especially when deployed on the cloud where their tail latency easily degrades. Replicating redundant requests has been shown to be an effective mechanism to defend application performance from high capacity variability. In this talk, I will present a PArtial REplication system, sPARE, that replicates and dispatches read-only workloads for distributed multi-tier web applications, determining replication factors per tier by explicitly considering the capacity variability at each tier. Further, to make redundancy strategy cost-effective, sPARE provides a dual scaler that dynamically adjusts the redundancy levels of virtual machines and requests by exploiting heterogeneity of VM performance and cost. Our results show that sPARE can improve the tail latency of MediaWiki and Solr by a factor of almost 2.7x and 2.9x, respectively. Moreover, sPARE can achieve a stringent target latency by using clones of wimpy VMs with a significant cost saving, compared to scaling brawny VMs that have better performance at a higher unit cost.

9.30 – 9.50 *Autonomous Strategies for Elastic Data Stream Processing in the Fog: from Centralized to Distributed Approaches*
Francesco Lo Presti, University of Rome “Tor Vergata”, Italy

In this talk we deal with the problem of elastic autonomous runtime management of Data Stream Processing (DSP) applications in a distributed environment as envisioned by the Fog computing paradigm. For this setting, we present a hierarchical and distributed architecture for the autonomous control of elastic DSP applications which attempts to overcome the limitation of centralized approaches. The proposed solution revolves around a two layered approach. At the lower level, distributed components issue requests for adapting the deployment of DSP operations as to adjust to changing working conditions. At the higher level, a per-application centralized component working on a broader time scale oversees the application behavior and grants reconfigurations to control the application performance while limiting the negative impact of application downtime. We implemented the proposed solution in our distributed Storm prototype and evaluated their behavior under realistic workload settings.

10.20 - 10.40 *AI: Beyond Machine Learning for Management*
Daniel Gillblad, Swedish Institute of Computer Science, Sweden

With the availability of massive computational resources, large data sets, and improvements in methods, Machine Learning has emerged as perhaps the most important component of smart, AI systems. Still, state-of-the-art ML methods come with severe restrictions in terms of data and computational efficiency, interpretability, and support for actual actionable knowledge. We will discuss these limitations, possible solutions, and how they apply to Cloud Management. Further, while Machine Learning is essential to AI, it only represents part of the field, with other components often equally important in creating an actual intelligent, autonomous application. Rather than applying ML to cloud management, we will argue for an "AI perspective" and provide examples of other areas of AI from which we can draw ideas and solutions.

10.40 - 12.00 Discussion 9: *Efficient Workload Modeling for future Software Defined Infrastructures*
Tomas Fredberg, Ericsson, Kista, Sweden

How can we describe (model) a cloud workload and its requirement on the infrastructure in a simple, yet efficient way that allow an automated and efficient composition, placement, scheduling and orchestration in a software defined infrastructure.

For further preparatory thoughts:
What metrics, tools and application workload modifications are needed to change the common physical description to a simpler and more resource optimized Support Level Objective into an SDI composition.

10.40 - 12.00 Discussion 10: *Managing the Trade-Off between Power and QoS of Applications: an Old Story in a New Heterogeneous World of Things*
Jakub Krzywda, Umeå University, Sweden and Monica Vitali, Politecnico di Milano, Italy

Managing applications involves also the management of their non functional aspects such as QoS and power consumption, which are usually in contrast with each other. This issue has been widely discussed in the past and several effective solutions have been proposed. In recent years, with the advent of the edge computing and microservices, the way in which applications are managed,

deployed, and distributed has significantly changed, adding several levels of complexity. A number of open issues have arisen:

- Enforcing the application power budgets on data center serves: how to ensure that the applications consume only the assigned power budget and that they are able to provide the acceptable performance? How to divide the total available power budget among the hosted applications to minimize the performance degradation?
- Modelling of application power-performance tradeoffs - how to capture the dependencies between the workload level, available power budget and the application performance?
- Power and QoS aware applications: how to make the application owner and data center operator aware of the power consumption and of the global QoS of an application composed of several microservices? How to account for data storage, usage, and movement when dealing with distributed applications?

All these issues will be considered during the discussion in order to discover interesting research challenges.

10.40 - 12.00 Discussion 11: *Making the Cloud Extremely Dependable*
Miguel Correia, INESC-ID/IST, Portugal

Clouds fail… a lot. News about outages are common, but other issues such as loss of files, hacking, and malicious insiders are also known to have happened recently. However, Dependability is a well established discipline, with several decades of scientific research and practical results. Its objective is do help producing dependable systems, i.e., systems that provide attributes such as integrity, availability, reliability, confidentiality, safety, and maintainability. The objective of this discussion session is to debate, and if possible produce a list of, guidelines for achieving dependable clouds. We will discuss mechanisms and processes for each dependability attribute and for each main architectural component. After a initial framework is produced, we will discuss how to go even further and increase even more the level of dependability. Alternatively, we may reach to the conclusion that the objective is unfeasible and that clouds will continue to fail… a lot.

10.40 - 12.00 Discussion 12: *The Grand Cluster Resource Utilization Challenge*
Marcel Blöcher, Technical University of Darmstadt, Germany

Cloud systems are large complex distributed systems characterized by heterogeneity, multi tenancy, geographical dispersion, and changes of workloads and resources. Various workloads share the same resources with different objectives and different constraints on locality and connectivity of their components. While nowadays resource provisioning of data center resources uses isolated compute container as resource shares, future cloud platforms additionally require insights into the container's connectivity beyond the border of a single data center. For effective resource scheduling, either applications thus have to learn more about the infrastructure, or the resource manager has to understand application semantics; Both options violate the separation of applications from the underlying infrastructure leading to the cluster resource utilization dilemma. In this session we will discuss issues, challenges, and limitations of all involved components including the data center scheduler and application framework level optimizer.

13.00 - 13.20 *Enos: a Holistic Framework for Conducting Scientific Evaluations of OpenStack*
Adrien Lebre, Inria / IMT-Atlantique, Nantes, France

By massively adopting OpenStack for operating small to large private and public clouds, the industry has made it one of the largest running software project. Driven by an incredibly vibrant community,

OpenStack has now overgrown the Linux kernel. However, with success comes an increased complexity; facing technical and scientific challenges, developers are in great difficulty when testing the impact of individual changes on the performance of such a large codebase, which will likely slow down the evolution of OpenStack. In the light of the difficulties the OpenStack community is facing, we claim that it is time for our scientific community to join the effort and get involved in the development and the evolution of OpenStack, as it has been once done for Linux. However, diving into complex software such as OpenStack is tedious: reliable tools are necessary to ease the efforts of our community and make science as collaborative as possible.
In this talk, I will give an overview of EnOS, an integrated framework that relies on container technologies for deploying and evaluating OpenStack on any testbed. EnOS allows researchers to easily express different configurations, enabling fine-grained investigations of OpenStack services. EnOS collects performance metrics at runtime and stores them for post-mortem analysis and sharing. The relevance of ENOS approach to reproducible research is illustrated by evaluating different OpenStack scenarios on the Grid'5000 and Chameleon testbeds.

13.20 - 13.35 *Checkpointing In-Memory Data Analytics Applications with Panda*
Bogdan Ghit, TU Delft, Netherlands

Providing fault-tolerance is of major importance for data analytics frameworks such as Hadoop and Spark, which are typically deployed in large clusters that are known to experience high failures rates. Unexpected events such as compute node failures are in particular an important challenge for in-memory data analytics frameworks, as the widely adopted approach to deal with them is to recompute work already done. Recomputing lost work, however, requires allocation of extra resource to re-execute tasks, thus increasing the job runtimes. To address this problem, we design a checkpointing system called Panda that is tailored to the intrinsic characteristics of data analytics frameworks. In particular, Panda employs fine-grained checkpointing at the level of task outputs and dynamically identifies tasks that are worthwhile to be checkpointed rather than be recomputed.

13.35 - 13.50 *Model-free QoS Control Using Cooperative Fuzzy Reinforcement Learning*
Olumuyiwa Ibidunmoye, Umeå University, Sweden

Designing efficient control mechanisms to meet strict performance requirements with respect to changing workload demands without sacrificing resource efficiency remains a challenge in cloud infrastructures. A popular approach to this problem is via fine-grained autoscaling mechanisms that rely on either simplistic threshold-based adaptation rules or sophisticated queueing/control-theoretic models to provision just the right amount of systems resources to cloud services. While it is difficult at design time to specify optimal thresholds or rules, it is even more challenging inferring precise models of workloads demands and execution context of the multitude of services in a cloud infrastructure. We propose CoFReL a model-free semi-supervised approach for guaranteeing performance QoS using cooperative fuzzy reinforcement learning for cloud services. CoFReL is a self-learning controller which incrementally learns and adjusts allocated compute capacity with limited supervision. CoFReL is currently being evaluated on real experimental testbed with realistic workloads. Preliminary results suggests that the approach is able, to a reasonable extent, meet performance SLA requirements and achieve resource efficiency without precise models of the execution environment.

Wednesday, June 14th

8.20 - 8.50 *ML for Datacenter Cooling Control - Progress and Challenges at Google*
Steve Webster, Google, US

Google has had some success using machine learning (ML) techniques to optimize energy usage in our datacenter cooling control systems (see https://deepmind.com/blog/deepmind-ai-reduces-google-data-centre-cooling-bill-40/). While our prior experience with non-ML realtime equipment control set an expectation that the ML aspects of the work would be most challenging, the actual effort has largely been in data collection and hygiene, failsafe control, and safe exploration.

We'll discuss lessons learned, progress so far, and opportunities for advancing ML training techniques in real-time control systems.

8.50 - 9.10 *Prediction Models for Service Responsiveness of Deep Neural Networks in a Cloud Setting*
Evgenia Smirni, College of William and Mary, US

Deep neural networks (DNNs) enable a host of artificial intelligence applications. These applications are supported by large DNN models running in serving mode often on a cloud computing infrastructure. Given the compute-intensive nature of large DNN models, a key challenge for DNN serving systems is to minimize user request response latencies. We show and model two important properties of DNN workloads that can allow for the use of queueing network models for predicting user request latencies: homogeneous request service demands and performance interference among requests running concurrently due to cache/memory contention. These properties motivate the design of a dynamic scheduling framework that is powered by an interference-aware queueing-based analytic model. The framework is evaluated in the context of an image classification service using several well known benchmarks. The results demonstrate its accurate latency prediction and its ability to adapt to changing load conditions, thanks to the fast deployment and accuracy of analytic queuing models. This work is in collaboration with Feng Yan of the University of Nevada at Reno, and Yuxiong He and Olatunji Ruwase of Microsoft Research.

9.10 - 9.30 *HOPS a Platform for Data Analytics in Data Centers*
Seif Haridi, Swedish Institute of Computer Science and KTH, Sweden

The talk with give an overview of the HOPS platform that is operational on SICS ICE the SICS data center. In particular we will cover HOPS-FS a scalable distributed file system that is compatible to HDFS. HOPS-FS scales the number of file operations up to 37 times that HDSF and the meta data from Giga bytes to Peta bytes. We will also describe the support of GPUs as resources in the resource manager HOPS-Yarn.

9.30 - 9.50 *Benchmarking Public Clouds With Cloud Workbench - Do You Really Know What You Are Paying For?*
Philipp Leitner, University of Zürich, Switzerland

Public IaaS cloud providers are notoriously vague when specifying what performance users can expect from their various instance types. What is even worse is that performance is often also highly unpredictable - the same instance type might be fast or slow, depending on what an instance's neighbours on the same rack or network segment are doing. This makes capacity planning in IaaS

clouds a challenging problem for users, often requiring substantial overpayment just to stay on the safe side. In this talk I will present Cloud Workbench, an open source tool that allows users to define and execute their own benchmarks against arbitrary cloud providers. I will also show some past example experiments and results to give you a flavour of what is possible with Cloud Workbench.

10.20 - 10.40 *Taking Elasticity to the Edge*
Srikumar Venugopal, IBM Ireland

One of the defining features of cloud computing is elasticity, or the ability to add or remove resources in response to workload demand. Commonly, cloud elasticity control mechanisms have relied on centralized approaches that use thresholds or user-defined functions to drive the provisioning of resources. Centralised architectures are complex, and may not be nimble enough to react to dynamic changes in workload. This presentation will present past research on autonomous, decentralized management of cloud elasticity. It will then discuss the unique challenges for elasticity posed by edge computing and how decentralization can help address these.

10.40 - 12.00 Discussion 13: *Causal Analytics for Automating Data Center Management*
James Kempf, Ericsson Research, US

While much attention today has been focused on machine intelligence (by which people usually mean neural networks) for data center management, a point that is often lost in the discussion is the need for determining the causal source for problems. Neural nets are very good at pattern recognition, but as examples in the image recognition domain have shown, if a few patches are missing in the image, the neural net can mistake a dog for a horse. Translating to the data center management domain, one would not want to mistake an intruder breaking into the data center for a perfectly harmless spike in user demand due to a new service. In this discussion, I'd like to have the group talk about causal analytic models for data center management, and how they can be used together with neural nets. I hope to jumpstart the discussion with a relatively primitive example of using Bayesian Graphical Networks for making autoscaling decisions. The hypothesis is that neural networks alone are not sufficient for automating data center management, causal analytics are also required, i.e. you need both.

10.40 - 12.00 Discussion 14: *Data Driven and Distributed Optimization in Cloud Management: Opportunities and Challenges*
Alessandro Papadopoulos, Mälardalen University, Sweden and Maria Prandini, Politecnico di Milano, Italy

Methods for data-driven and distributed optimization in a multi-agent system affected by uncertainty have been introduced recently to address resource allocation problems in energy systems.
These methods appear to have potential for application to resource allocation problems in cloud management where "data" are requests of computational resources (CPU, memory, etc).
The goal of this discussion session is to work out opportunities and challenges for data-driven and distributed optimization in cloud management, building on the complementary expertises of participants and organizers.

10.40 - 12.00 Discussion 15: *Command & Conquer: Control and Assurance in Distributed Systems*
Joe Butler and Thijs Metsch, Intel labs, Leixlip, Ireland

Scale and distribution of resources are becoming key challenges as we shift from more centralized cloud computing deployments towards the edge in pursuit of satisfying key requirements such as

lowering latency. This requirement is particularly relevant for emerging use cases in transportation (autonomous driving), industry 4.0, and similar CPS-enabled systems.

In this session we want to discuss how we can a) enable control in such distributed systems and b) ensure service delivery through these mechanisms. What techniques offer promise? Entities expressing their utility towards a global objective at scale? … ?

10.40 - 12.00 Discussion 16: *Clouds & 5G: Interplay, Opportunities and Challenges*
Dimosthenis Kyriazis, University of Piraeus, Greece and Luis Tomas, Red Hat, Spain

Cloud computing environments have reached a level of maturity, being considered as a "utility" and exploited by a variety of applications in different contexts. The current shift to the edge as well as the increasing data-related needs of applications (going beyond centralized data processing to multiple sources and cross-stream processing) pose specific requirements to the underlying network infrastructures. The goal of the discussion in this session is to identify how the interplay between clouds and 5G environments creates opportunities for both domains, while raising challenges that need to be addressed in order for both of them to reach their full potential and exploit the outcomes of the aforementioned interplay.

13.00 - 13.20 *The Evaporating Cloud*
Guillaume Pierre, University Rennes 1, France

Cloud computing data centers are composed of very powerful computing nodes connected by reliable backbone networks. However, these resources are concentrated in a small number of a data centers. The latency between an end user and the closest available cloud data center comes in the range of 20-150 ms. A number of latency-sensitive applications (e.g., augmented reality) require extremely low end-to-end latencies and therefore cannot make use of traditional cloud platforms. Fog computing therefore aims to complement traditional cloud infrastructures with additional resources located extremely close to the user, within a couple of network hops. This requires one to distribute machines in a very large number of geographical locations so computation capacity is always available in immediate proximity of any end user. In this presentation I will discuss the architectural challenges one needs to face when designing the next-generation fog computing architectures.

13.20 - 13.40 *Event-Driven Bandwidth Allocation to a Network of Cameras with Formal Guarantees*
Martina Maggio, Lund University, Sweden

Modern computing systems are often formed by multiple components that interact with each other through the use of shared resources (e.g., CPU, network bandwidth, storage). We consider system that consists of a network of self-adaptive cameras that share a communication channel, transmitting streams of frames to a central node. The cameras can modify a quality parameter to adapt the amount of information encoded and to affect their bandwidth requirements and usage. A critical design choice for such a system is scheduling channel access, i.e., how to determine the amount of channel capacity that should be used by each of the cameras at any point in time. Two main issues have to be considered for the choice of a bandwidth allocation scheme: (i) camera adaptation and network access scheduling may interfere with one another, (ii) bandwidth distribution should be triggered only when necessary, to limit additional overhead. The talk will present a formally verified (using Model Checking) event-triggered adaptation scheme for bandwidth allocation, designed to minimize additional overhead in the network.

13.40 - 13.55 *An Efficient Communication Aware Heuristic for Multiple Cloud Application Placement*
Pedro Silva, Inria, ENS de Lyon, France

To deploy a distributed application on the cloud, cost, resource and communication constraints have to be considered to select the most suitable Virtual Machines (VMs), from private and public cloud providers. This process becomes very complex in large scale scenarios and, as this problem is NP-Hard, its automation must take scalability into consideration. In this presentation, we propose a heuristic able to calculate placements for distributed component-based applications on possibly multiple clouds with the objective of minimizing VM renting costs while satisfying applications' resource and communication constraints. Our evaluation results indicate that the proposed heuristic is able to compute good quality solutions much faster than state of the art placement approaches, namely exact algorithms and meta-heuristics.

13.55 - 14.10 *Dynamic Control of NFV Forwarding Graphs with End-to-end Constraints*
Victor Millnert, Lund University, Sweden

There is a strong industrial drive to use cloud computing technologies and concepts for providing timing sensitive services in the networking domain since it would provide the means to share the physical resources among multiple users and thus increase the elasticity and reduce the costs. In this talk, we show a mathematical model for user-stateless virtual network functions forming a forwarding graph. The model captures uncertainties of the performance of these virtual resources as well as the time-overhead needed to instantiate them. The model is used to derive a service controller for horizontal scaling of the virtual resources as well as an admission controller that guarantees that packets exiting the forwarding graph meet their end-to-end deadline.

14.10 - 14.25 *Auto-scaling Workflows: A Comparative Study*
Ahmed Ali-Eldin, Umeå University, Sweden

Cloud applications are very diverse in properties ranging from request-response workloads to task based workloads. To manage these applications, many autoscaling policies have been proposed in the past decade controlling when and how to provision resources to a cloud application utilizing cloud elasticity features. However, in prior work, when a new policy is proposed, it is seldom compared to the state-of-the-art, and is often compared only to static provisioning using a predefined QoS target. This reduces the ability of cloud customers and of cloud operators to choose and deploy an autoscaling policy. In our work, we conduct an experimental performance evaluation of autoscaling policies, using workflows, a commonly used formalism for automating resource management for applications with well-defined yet complex structure. We present a detailed comparative study of general state-of-the-art autoscaling policies, along with two new workflow-specific policies. To understand the performance differences between the 7 policies, we conduct various forms of pairwise and group comparisons. We report both individual and aggregated metrics. Our results highlight the trade-offs between the suggested policies, and thus enable a better understanding of the current state-of-the-art.