Program

13th Cloud Control Workshop

Skåvsjöholm, Åkersberga, Sweden
June 13 – 15, 2018

Practical information

Hotel contact:
Skåvsjöholm Konferens & Möten
Address: Skåvsjöholmsv. 80, 184 94 Åkersberga
e-mail: info@skavsjoholm.se
Telephone: +46-8-540 267 00

Hotel reception opening hours:
07:30-21:00
Dial 9 to call from your room

Check-in and -out
Hotel room check-in from 14:00. Check-out until 10:00. (For check-out before 7.30, contact the reception the day before.)

Payment
Rooms, meals, and some drinks are covered by the workshop registration fee. Please ensure to pay at check-out any additional costs put on your room, like telephone, additional drinks, etc.

Cafeteria
Cafeteria on entrance floor with coffee, the, juice, fruit, and biscuits available free of charge during 7.30-20.00.

Luggage room
The conference room Skarpö on entrance floor is dedicated as our luggage room on arrival.

Relax facilities and other leisure activities
The relax facilities includes a sauna and outdoor whirlpool and is by default open 17.00-22.00 (from 16:00 on June 14). For use at other hours, tell the reception an hour in advance.
Equipment for various social games (e.g., boule, crocket, kubb) are available by the relax facilities (on shelfs near the bar)
Bicycles: ask at the reception
Canoes: ask at the reception – do not forget the life wests!

Bar
The main bar is located at the ground floor (near the relax facilities) and is open 17.00-24.00.
There is also a lobby bar by the entrance.

Breakfast
Breakfast is served in the main dining room at 07.00-09.00

Coffee breaks
Coffee during coffee breaks is served in the main dining room

Meals and drinks
Lunches are served in the main dining room
Wednesday 2-course dinner is served in the main dining room. Dessert is from a buffet and can be brought to outside or to the lounge area.
Thursday barbecue buffet dinner is served outdoors, by the sea.
At workshop registration, all participants receive a few drink tickets to be used for the Poster Reception and dinners.

Internet
WiFi throughout the premises
In main building: "SKAVAN" – No password
In Villa Skåvsjöholm: ”VILLAN” - No password
In Husarö: "HUSARO" Password: welcometoskavan

Weather
There has been great summer weather for weeks now, but forecasts during the workshop is less stable weather and slightly chillier. Bring some slightly warm clothing for the outdoor barbeque dinner on Thursday night.

In case everything else fails
Erik’s phone number is +46 70 315 3928. Anticipate limited ability to answer around bus departure on Wednesday.

Wednesday June 13th

8:40 Gathering at bus outside Arlanda airport, Terminal 4

9:00 Bus departure

9:40 Coffee, registration, and luggage storage (in the room Skarpö on entrance floor)

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

10:10 *Workshop Introduction*
Erik Elmroth, Umeå University, Sweden

10:25 **Keynote: *Designing Systems and Applications for Transient Computing***
Prashant Shenoy, University of Massachusetts Amherst, USA

11:10 *Challenges in Scaling up Datacenter ML Control*
Steve Webster, Google, USA

11:30 *Dynamically Scaling Container-based Microservices Architectures*
Arno Jacobsen, MSRG TU Munich, Germany / University of Toronto, Canada

11:50 *Differential Calculus for the Cloud: A/B Testing and How to Speed up the Software Development Cycle*
Simon Tuffs, Palo Alto Networks, USA

12:10 *Predicting Failures in Cloud Systems*
Mauro Pezzè, Università della Svizzera italiana, Switzerland

12:30 **Lunch**

Session Chair: Azimeh Sefidcon, Ericsson Research, Sweden

13:45 *Cloud Management in Twister2: A High-Performance Big Data Programming Environment*
Geoffrey Fox, Indiana University Bloomington, USA

14:05 *"Information Models" and Extracting more Value from Volatile Resources*
Andrew A. Chien, University of Chicago, USA

14:25

Discussion 1 (Resarö)	Discussion 2 (Rindö)	Discussion 3 (Fåglarö)	Discussion 4 (Huvudskär)	Discussion 5 (Svavelsö)
Autonomous datacenters Simon Tuffs, Palo Alto Networks, US & Karl-Erik Årzén, Lund University, Sweden	*Tracing distributed systems in the Cloud* Rodrigo Fonseca, Brown University, US	*Performance boosting for cloud network algorithms – Can artificial intelligence help?* Andreas Blenk & Johannes Zerwas, TU Munich, Germany	*Principles for testing and experimentation* Alexandru Iosup, Vrije Universiteit Amster-dam, The Netherlands & Alessandro Papado-poulos, Mälardalen University, Sweden	*A rigorous Cloud: is there a case for formal methods?* Wolfgang John, Ericsson Research, Sweden & Stefan Schmid, University of Vienna, Austria

15:45 **Group Photo Shoot** (Mikael Hansson, Umeå University, Sweden)

15.55 **Coffee & Hotel Room Check-in**

Session Chair: Frederic Desprez, INRIA, France

16.25 *Programming the Cloud Computer using Serverless Compositions*
Paul C. Castro, IBM, USA

16:55 ***Poster Reception with Refreshments*** (Svavelsö)

1. *Network Algorithms Ex Machina: Towards Self-Driving Communication Networks*, Andreas Blenk, TU Munich, Germany
2. *Effect of Data imputation on Classification Performance: Case study of Datacenter Trouble Tickets,* Monowar Bhuyan, Umeå University, Sweden
3. *Kairos: Preemptive Data Center Scheduler Without Runtime Estimates*, Pamela Delgado, EPFL, Switzerland
4. *ALPACA: Application Performance Aware Server Power Capping*, Jakub Krzywda, Umeå University, Sweden
5. *Privacy-Preserving Data Processing on the Network Edge*, Lars Larsson, Umeå University, Sweden
6. *Distributed Placement for Flow-based Applications in Edge Clouds*, Amardeep Mehta, Umeå University, Sweden
7. *Deadline Disco*, Victor Millnert, Lund University, Sweden
8. *Location-aware Workload Prediction for Edge Data Centers*, Chanh Nguyen, Umeå University, Sweden
9. *Cloud Application Predictability through Integrated Load-Balancing and Service Time Control*, Tommi Nylander & Marcus Thelander Andrén, Lund University, Sweden
10. *Online Anomaly Detection and Real-Time Causality Recommendation*, Mohamad Rezaei, KTH, Sweden
11. *Server State Estimation Using Event-Based Particle Filtering*, Johan Ruuskanen, Lund University, Sweden
12. *Mission Critical Cloud*, Per Skarin, Ericsson/Lund University, Sweden
13. *Hybrid Resource Management for HPC and Data Intensive Workloads*, Abel Souza, Umeå University, Sweden
14. *Resource Management in Large Scale Data Centres Using Distributed Optimisation Algorithms*, Javad Zarrin, University of Cambridge, UK
15. *Towards Flexible and Adaptive Networks: From Topology Design to Data Center Resource Management*, Johannes Zerwas, TU Munich, Germany

19:00 **Dinner**

Thursday June 14th

7:00 **Breakfast**

Session Chair: Joe Butler, Intel Labs, Ireland

8:15 ***Keynote: Deep-Learning-as-a-Service for IoT Systems***
Tarek Abdelzaher, University of Illinois, USA

9:00 *Interference-aware Resource Management in Cloud Environments*
Anshul Gandhi, Stony Brook University, USA

9:20 *Never Late Again! Job-Level Deadline SLOs in YARN*
Subramaniam Venkatraman Krishnan, Microsoft, USA

9:40 **Coffee**

Session Chair: Maria Kihl, Lund University, Sweden

10:10 *Privacy Preservation for WiFi-tracking Data*
Maarten van Steen, University of Twente, The Netherlands

10:30 *Synchronization Models for Scalable Machine Learning in the Cloud*
Lixin Gao, University of Massachusetts Amherst, USA

10:50 *Navigating Heterogeneity and Uncertainty in Private Cloud Scheduling*
Bhuvan Urgaonkar, The Pennsylvania State University, USA

11:10

Discussion 6 (Resarö)	Discussion 7 (Rindö)	Discussion 8 (Storskär)	Discussion 9 (Huvudskär)	Discussion 10 (Svavelsö)
Real world challenges in cloud operations Narayan Desai, Google, US & Johan Eker, Ericsson Research / Lund University, Sweden	*Operational analytics at cloud grade* Eran Raichstein, IBM Research, Israel	*Making big data processing just work* Bogdan Ghit, Databricks, The Netherlands & Cristian Klein, Umeå University, Sweden	*A Tale of Distribution: Continuous and Event-based Software Engineering Strategies for Cloud Management* Martina Maggio, Lund University, Sweden	*Distributed intelligence: How can Artificial Intelligence help with the challenges of managing and controlling dynamic Fog/Edge compute platforms?* Joe Butler & Thijs Metsch, Intel Labs, Ireland

12:30 **Lunch**

Session Chair: Rolf Stadler, KTH, Sweden

13:45 **Keynote*: Building the Warehouse Scale Computer***
John Wilkes, Google, USA

14.30

Discussion 11 (Resarö)	Discussion 12 (Rindö)	Discussion 13 (Storskär)	Discussion14(Huvudskär)	Discussion 15 (Svavelsö)
Cloud Computing for an AI First future Geoffrey Fox, Indiana University Bloomington, USA	*Hybrid cloud orchestration, can we make it simple?* Gal Hammer, Red Hat, Israel & Luis Tomas, Red Hat, Spain	*Software Defined Infrastructure – Opportunities and challenges* Mats Jonsson, Saab AB, Sweden	*Edge Clouds: Limitations and Applications* Ahmed Ali-Eldin, Umeå University, Sweden / Umass, Amherst, US & Farah Ait Salaht, INRIA-Grenoble, France	*Practices and challenges for building cloud-native applications with FaaS,* Philipp Leitner & Joel Scheuner, Chalmers / University of Gothenburg, Sweden

15:45 **Coffee**

16:15 Social outdoor activities
Equipment and facilities available at no cost:

- Equipment for various social games (e.g., boule, crocket, kubb) are available by the relax facilities (on shelfs near the bar)
- Outdoor whirlpool and indoor sauna is available by the relax facilities
- Bicycles: ask at the reception
- Canoes: ask at the reception – do not forget the life wests!

19:00 **Barbeque Dinner**

Friday June 15th

7:00 Breakfast

Session Chair: Seif Haridi, KTH / SICS, Sweden

8:15 ***Keynote**: **Shared Memory and Disaggregated Memory in Virtualized Clouds***
Ling Liu, Georgia Institute of Technology, USA

9:00 *Harvesting Randomness to Optimize Distributed Systems*
Siddharta Sen, Microsoft, USA

9:20 *Networking as a First-Class Cloud Resource*
Rodrigo Fonseca, Brown University, USA

9:40 Coffee & Hotel Room Check-out (before 10.00)

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

10.10 *WIP: Exploiting Heterogeneity to Minimize Tail Latency*
Timothy Zhu, The Pennsylvania State University, USA

10:30 *Decentralized Self-Adaptation for Elastic Data Stream Processing*
Valeria Cardellini, University of Rome Tor Vergata, USA

10:50 *Back to the Future: Towards Timely Bounded Cloud Communication*
Patrick Eugster, University of Lugano, Switzerland

11:10

Discussion 16 (Resarö)	Discussion 17 (Rindö)	Discussion 18 (Storskär)	Discussion19(Huvudskär)	Discussion 20 (Svavelsö)
Cloud versus in-network-processing for latency-critical industrial control operations Anders Robertsson, Lund University, Sweden & Klaus Wehrle, RWTH Aachen University, Germany	*Customised Data as a Service for multiple and conflicting data intensive applications in cloud/edge environments* Monica Vitali, Politecnico di Milano, Italy	*Managing Container Clouds: Challenges and Opportunities* Ali Kanso, IBM, US	*Data plane program-mability for flexible cloud networking* Andreas Kassler, Karlstad University, Sweden	*Single System Image and distributed OS: Has the time come?* Daniel Turull & Mina Sedaghat, Ericsson Research, Sweden

12.30 Lunch

Session Chair: Guillaume Pierre, Université de Rennes 1, France

13.45 *Reducing the Cost of Exploring Deployment Configurations in the Cloud*
Luis Rodrigues, INESC-ID, IST, Ulisboa, Portugal

14:05 *Massivizing Computer Systems: a Vision to Understand, Design, and Engineer Computer Ecosystems through and beyond Modern Distributed Systems*
Alexandru Iosup, Vrije Universiteit Amsterdam, The Netherlands

14:25 *Experimental Evaluation of Amazon EC2 Spot Instances*
Thomas Fahringer, University of Innsbruck, Austria

14:45 *Closing*
Erik Elmroth, Umeå University, Sweden

15:00 Coffee

15:30 Bus departure

16:15 Bus arrival at Arlanda airport, Terminal 4

Presentations, Posters, and Discussion Sessions in order of Appearance
13th Cloud Control Workshop
June 13 – 15, 2018

Wednesday, June 13th

10.10 *Workshop Introduction*
Erik Elmroth, Umeå University, Sweden

We will start by a short welcome and to give a brief historic flashback on how the project meetings in a project named Cloud Control (2nd largest project grant ever awarded by The Swedish Research Council) grow into an international workshop with no connections to that finished project other than the name (and the fact that Umeå and Lund Universities are still over-represented among participants). The main driver behind that transformation – *To run meetings whose forms dynamically adjust to make them as useful as possible for leading researchers in our field* – is still what drives the development of the series. As the main attraction of the workshop is the 107 participants themselves, we will also reflect on who they are. As for statistics, they represent organizations in 15 countries – 50% from Sweden but also 20% from the US and at least three people from each of Switzerland, Germany, France, Norway, Italy, and the Netherlands. Industries like Ericsson Research, Google, IBM Research, Microsoft Research, Intel Labs, Red Hat, Palo Alto Networks, Databricks, and more are represented by 28 participants. From Universities, there are 29 Full Professors as well as 14 other Faculty members, 21 PhD students, and many additional researchers and postdocs. If excluding the over-represented teams from Umeå and Lund Universities, over 60% of the participants are attending their first Cloud Control Workshop.

10:25 **Keynote: *Designing Systems and Applications for Transient Computing***
Prashant Shenoy, University of Massachusetts Amherst, USA

Traditional distributed systems are built under the assumption that system resources will be available for use by applications unless there is a failure. Transient computing is a new phenomena that challenges this assumption by allowing system resources to become unavailable at any time. Transiency arises in many domains such as cloud computing--in the form of revocable spot servers--and in data centers that rely on variable electricity prices or intermittent renewable sources of energy. Transiency is inherently different from fault tolerance since resources do not fail, rather they become temporarily unavailable, and traditional fault tolerance mechanisms are not suitable for handling transient resource unavailability.

In this talk, I will discuss how systems and applications need to be rethought to run on transient computing systems. I will first describe a system called Yank that uses a new bounded-time virtual machine migration mechanism to handle transiency at a system level while being transparent to applications. I will then discuss how modern distributed applications can be made transiency-aware and present a Spark-variant called Flint that we have developed to exploit transient cloud computing. I will end with open research questions in this area and directions for future work.

11:10 *Challenges in Scaling up Datacenter ML Control*
Steve Webster, Google, USA

Following on Google's early success with ML control of our datacenter cooling systems has been a series of learning experiences and new challenges. In this brief talk, I'll discuss how we've discovered the hard parts of ML control of large-scale systems typically aren't the ML at all, but rather the data quality, human operators, and configuration.

11:30 *Dynamically Scaling Container-based Microservices Architectures*
Arno Jacobsen, MSRG TU Munich, Germany / University of Toronto, Canada

Frequently changing requirements demand for increasingly innovative software architectural styles that accommodate application scalability and development flexibility. Recently, microservices architectures have garnered the attention of many organizations, providing higher levels of scalability, availability, and fault isolation. Many organizations choose to host their microservices architectures in cloud data centres to offset costs. Incidentally, data centres become over-encumbered during peak usage hours and underutilized during off-peak hours. In this talk, we sketch two new autoscaling algorithms and benchmarked them against Google's popular Kubernetes autoscaling algorithm. Results indicated up to 1.49x speedups in response times, 10 times less failed requests, and 26% increase in resource efficiencies.

11:50 *Differential Calculus for the Cloud: A/B Testing and How to Speed up the Software Development Cycle*
Simon Tuffs, Palo Alto Networks, USA

One of the enduring challenges in software development is delivering system upgrades in an efficient manner, without affecting availability. The introduction of micro-systems architectures in cloud computing allows a conceptual shift from a develop-test-deploy-operate lifecycle to one driven by data. This paper examines the use of a "differential equation" to answer the question: "how do we know our new code will work in production, correctly and better than our existing code, before we commit to it?". Infrastructure and analytics are used to close feedback loops around and between various lifecycle phases. There are demonstrable benefits from this "A/B" approach which exceed the cost of the required infrastructure and processes. The A/B approach raises some interesting statistical and information theoretical questions such as "how rich is the data stream relative to the complexity of the code under development?", and "how confident are we that the system is going to work in production?". An example taken from the Palo Alto Networks Aperture Machine-Learning Classifier project will be used as an illustration.

12:10 *Predicting Failures in Cloud Systems*
Mauro Pezzè, Università della Svizzera italiana, Switzerland

Runtime failures are unavoidable in the cloud, and predicting failures before their occurrence is a viable solution to either prevent the occurrence of failures in the cloud or mitigate their effects. We present an alternative lightweight and precise approach to predict failures and locate the corresponding faults in cloud systems. The approach blends anomaly-based statistical analysis techniques to identify anomalous behaviors and signature-based semi-supervised machine learning techniques to precisely predict failures. The talk frames the problem, presents the approach, discusses experimental result that confirm the precision and efficiency of the approach, and illustrates constraints and requirements that derive from an ongoing per-industrial application of the approach.

13:45 *Cloud Management in Twister2: A High-Performance Big Data Programming Environment*
Geoffrey Fox, Indiana University Bloomington, USA

We discuss Twister2 which consists of a set of middleware components to support batch or streaming data capabilities familiar from Apache Hadoop, Spark, Heron and Flink but with high performance. Twister2 covers bulk synchronous and data flow communication; task management as in Mesos, Yarn and Kubernetes; dataflow graph execution models; launching of the Harp-DAAL machine-learning library; streaming and repository data access interfaces, in-memory databases and fault tolerance at dataflow nodes. We discuss how cloud management functions could be added.

14:05 *"Information Models" and Extracting more Value from Volatile Resources*
Andrew A. Chien, University of Chicago, USA

Cloud providers sell unreliable or ``volatile" resources that are unused by foreground (reserved/high priority) workloads. The value users can extract from these resources depends on the (i) volatile resource management algorithm, and (ii) information provided to users about the resources. We evaluate volatile resource management and several information models (MTTR, limited statistics, Full distribution, and Oracle) using commercial cloud resource traces from 608 Amazon EC2 instance pools. The results show volatile resource management algorithms can increase user value by 30 to 45\%. Slightly richer information models (90pctile) combined with LIFO and LIFO-pools volatile resource management increase user value by as much as 10-fold.

14:25 Discussion 1 (Resarö)
Autonomous datacenters
Simon Tuffs, Palo Alto Networks, US & Karl-Erik Årzén, Lund University, Sweden

There is currently a substantial research effort devoted to cloud autonomy or cloud control. However, most of these efforts are aimed at improved resource management and performance or reduced environmental impact, utilizing techniques such as control, optimization, and analytics. However, an area where much less effort is spent from academia is datacenter automation, i.e, to what extent it is possible to automate the tasks of the datacenter operators, using, e.g., work flow automation or expert system techniques. However, in order for this it is essential to know what the operators spend their time on, what the tasks consists of, which decisions are made, what is the nature of the reasoning performed, what data and information that is used, and what the prospect are for automating these tasks. Are there any recent investigations of this?

14:25 Discussion 2 (Rindö)
Tracing distributed systems in the Cloud
Rodrigo Fonseca, Brown University, US

End-to-end (causal) tracing of distributed systems brings back some of the visibility that is lost when systems become collections of independent components, running on different machines, and written by different people. However, they are effective when most or all of the components agree to propagate and generate tracing information. Running on the cloud (or clouds) both increases the need, and utility, of distributed tracing, but it also makes it more difficult. What are the challenges going forward, to make this more pervasive? Interesting issues include standardization, providers' willingness to expose execution information, and incentives for such integration of tracing. Further challenges also include how to make the best use of such rich data to actually understand how systems operate, and how they sometimes fail.

14:25 Discussion 3 (Fåglarö)
Performance boosting for cloud network algorithms - Can artificial intelligence help?
Andreas Blenk & Johannes Zerwas, TU Munich, Germany

Networks are becoming more and more complex; humans struggle in designing new networking algorithms that satisfy in particular users and applications with their continuously increasing requirements. While over the last decades various technologies emerged that facilitate a more flexible resource management, like Cloud Computing and Network Virtualization, one problem still remains when managing networks: the hardness of the to-be-solved networking problems. So, can Artificial Intelligence (AI) help? One front where AI can assist is speed up algorithms, as recently shown in some research works. Another big potential of AI (and Machine Learning) is that it can be data-driven: considering algorithm data enables new solutions that help to advance existing traditional optimization procedures. However, AI may also lead to a loss of control. We would like foster the discussion of applying the data-driven paradigm to networking algorithms. Are ML-based algorithms and network designs already feasible? Or do we just inherit the same obstacles like curse of dimensionality as in traditional machine learning areas?

14:25 Discussion 4 (Huvudskär)
Principles for testing and experimentation
Alexandru Iosup, Vrije Universiteit Amsterdam, The Netherlands & Alessandro Papadopoulos, Mälardalen University, Sweden

The rapid adoption and the diversification of cloud computing technology exacerbate the importance of sound experimental methodology for this domain. This discussion deals with the concept of reproducibility in cloud computing experiments, how to measure and report performance in the cloud, and how well the cloud research community is already doing it.

14:25 Discussion 5 (Svavelsö)
A rigorous Cloud: is there a case for formal methods?
Wolfgang John, Ericsson Research, Sweden & Stefan Schmid, University of Vienna, Austria

While formal methods are currently conquering the networking community by storm, and researchers around the globe are developing tools for the automated testing and synthesis of provably correct communication networks, we have so far seen much less impact of such rigorous methods on the Cloud. In this session, we want to discuss possible application domains as well as the potential and limitations of formal methods in the context of Cloud. To get the discussion going, we will briefly present some applications of rigorous approaches in networking, such as a self-stabilizing control plane for software-defined networks, inspired by Radia Perlman's vision of a self-stabilizing Internet. We will then collaboratively explore similar opportunities, e.g., for Cloud control planes. The goal of this session is to identify interesting research directions, potential opportunities for collaborations, and perhaps even develop a small taxonomy of formal methods for the Cloud.

16.25 *Programming the Cloud Computer using Serverless Compositions*
Paul C. Castro, IBM, USA

Serverless computing is fostering a new paradigm in cloud application development and developer productivity. Its primary appeal is rooted in three key benefits: automatic elastic scaling, a pay-as-you go billing model at milliseconds intervals, and shifting the infrastructure operations and security entirely to the cloud vendor. Functions-as-a-Service (FaaS) is currently the canonical example of serverless. Most FaaS computing platforms assume functions are stateless and complete execution on the order of minutes. However, most real applications require the management of some state and may be long running. We present Composer, our work on a serverless orchestration layer that addresses these concerns. Composer is a unified programming model that exposes the three tenants of distributed and reactive programming for the serverless cloud: fine-grained computation with functions, explicit state for application data, and control flow for composing functions and services. With these programming features, it is possible to build enterprise scale applications, that are cloud native by construction. We present the Composer developer experience using an experimental, modern development environment we created called Cloud Shell, which integrates streamlined commands with interactive visualizations.

16:55 ***Poster Reception with Refreshments*** (Svavelsö)

Poster 1: Network Algorithms Ex Machina: Towards Self-Driving Communication Networks
Andreas Blenk, TU Munich, Germany

With networks becoming more and more complex, humans struggle in designing new solutions, e.g, network protocols, topology designs, or algorithms, that satisfy in particular users and applications with their continuously increasing requirements. While over the last decades various technologies emerged that facilitate a more flexible resource management, like Cloud Computing and Network Virtualization, one problem raised: the large amount of various decision possibilities, protocol configurations, and the hardness of the to-be-solved networking problems. Hence, it became hard to adapt today's communication networks. This limitation of today's communication networks result in poor user experience, network outages, high latency, interference, etc. We initiate the exploration of a new networking paradigm which we call Networks Ex Machina: networks that are optimized or even self-optimized towards their demand, usage, and context, in an automated manner. We see a clear demand for new paradigms that can handle arising amount of data and with it the exploding optimization space when facing networking problems. As a use case for self-driving networks, this poster presents our first studies of intelligence and machine- and data-driven networking algorithms for communication networks. Using machine learning and neural computation, we demonstrated that network algorithms can be learned and sped up: We have gained first experiences with learning existing network resource allocation algorithms, such as virtual network embedding and facility location.

Poster 2: Effect of Data imputation on Classification Performance: Case study of Datacenter Trouble Tickets,
Monowar Bhuyan, Umeå University, Sweden

Missing values are unavoidable in real-world performance datasets, especially when dealing with operational data from large cloud datacenters. The objective of this study is to investigate the effect of data imputation techniques on predictive power under high missingness. We analyze labelled trouble tickets data obtained from over 10K compute nodes across five commercial datacenters over a period of one year with over 70% missing values along critical features. To meet the study objective, we evaluate the ability of a variety of classifiers to automatically categorize tickets into one of six probable fault classes after filling missing data using single-local, single-global, as well as multiple-global imputation schemes. Initial results suggest that local imputation results in overly optimistic models while global imputation results in pessimistic models due to the high degree of missingness.

Poster 3: Kairos: Preemptive Data Center Scheduler Without Runtime Estimates
Pamela Delgado, EPFL, Switzerland

The vast majority of data center schedulers use job runtime estimates to improve the quality of their scheduling decisions. Knowledge about the runtimes allows the schedulers, among other things, to achieve better load balance and to avoid head-of-line blocking. Obtaining accurate runtime estimates is, however, far from trivial, and erroneous estimates may lead to sub-optimal scheduling decisions. Techniques to mitigate the effect of inaccurate estimates have shown some success, but the fundamental problem remains. This paper presents Kairos, a novel data center scheduler that assumes no prior information on job runtimes. Kairos introduces a distributed approximation of the Least Attained Service (LAS) scheduling policy. Kairos consists of a centralized scheduler and a per-node scheduler. The per-node schedulers implement LAS for tasks on their node, using preemption as necessary to avoid head-of-line blocking. The centralized scheduler distributes tasks among nodes in a manner that balances the load and imposes on each node a workload in which LAS provides favorable performance. We have implemented Kairos in YARN. We compare its performance against YARN FIFO scheduler and Big-C, an open source state-of-the-art YARN-based scheduler that also uses preemption. We show that Kairos reduces the median job completion time by 37% (resp. 73%) and the 99-th percentile by 57% (resp. 30%), with respect to Big-C and FIFO. We evaluate Kairos at scale by implementing it in the Eagle simulator and comparing its performance against Eagle. Kairos improves the 99th percentile of short job completion times by up to 55% and 85% the for the Google and Yahoo traces respectively.

Poster 4: ALPACA: Application Performance Aware Server Power Capping
Jakub Krzywda, Umeå University, Sweden

Server power capping limits the power consumption of a server to not exceed a specific power budget. This allows data center operators to reduce the peak power consumption at the cost of performance degradation of hosted applications. Previous work on server power capping rarely considers Quality-of-Service (QoS) requirements of consolidated services when enforcing the power budget. Therefore, we introduce ALPACA, a framework to reduce QoS violations and overall application performance degradation for consolidated services.

Our framework reduces unnecessary high power consumption when there is no performance gain, and divides the power among the running services in a way that reduces the overall QoS degradation when the power is scarce. ALPACA consists of the application power-performance modeling component that captures the relation between the power budget and application performance, an optimiser that chooses the best combination of power budget levels for each application, and a controller that enforces the optimal power budgets. We extensively evaluate the framework, performing over 300 testbed experiment runs and over 1000 simulation runs using four different applications, namely, a

full replica of the German Wikipedia, Sock Shop, SysBench, and CloudSuite Web Search benchmark under multiple realistic scenarios. Our experiments show that ALPACA reduces the operational costs of QoS penalties and electricity by up to 40% compared to a non-optimized system.
We are currently working on enhancing application power-performance models with hardware performance counters in order to provide an on-line performance feedback during the application runtime, even when the application does not expose own performance metrics. Our early results suggest that enhancing application power-performance models with hardware performance counters may enable more robust and holistic control of distributed applications hosted in cloud data centers.

Poster 5: Privacy-Preserving Data Processing on the Network Edge
Lars Larsson, Umeå University, Sweden

Privacy expectations or requirements make processing of some sensitive datasets impossible in current cloud offerings. We argue that neither typical Software-as-a-Service platforms nor traditional software licensing provide adequate ways to address this concern. Here, we present our serverless edge-cloud platform vision, in which statelessness and privacy-preservation is enforced by forbidding processes' access to the general network and permanent storage. To ensure that economic incentives of all parties are handled correctly, we introduce novel fine-grained micro-licences and processing proofs. The main purpose of the former is to guarantee that software vendors will receive payment for data processing, and the main purpose of the latter ensures customers that processing has finished correctly. Our research also clearly outlines future work we aim to undertake in this field.

Poster 6: Distributed Placement for Flow-based Applications in Edge Clouds
Amardeep Mehta, Umeå University, Sweden

Mobile Edge Clouds (MECs) with 5G are expected to create new opportunities for latency-critical applications in domains such as intelligent transportation systems, process automation, and smart grids. However, it is not clear yet how one can cost-efficiently deploy and manage a large number of such applications given the heterogeneity of devices, application performance requirements, and workloads.
This work explores cost and performance dynamics for IoT applications, and proposes distributed algorithms for automatic deployment of IoT applications in heterogeneous environments. Placement algorithms were evaluated with respect to metrics including number of required runtimes, applications' slowdown, and the number of iterations used to place an application. We are currently working on placement in Mobile Edge Clouds with capacity constraints where the objective is to minimize both operational cost and penalty due to QoS violations. Applications are classified at different priority levels based on their criticality. A scheduler at a runtime dynamically allocates resources to actors in order to minimize the QoS violations based on the applications' priority. It ensures fairness for applications at equal priority level. The placement is reevaluated periodically by a centralized Autonomic Runtime Allocator in order to minimize both Opex and penalty cost.

Poster 7: Deadline Disco
Victor Millnert, Lund University, Sweden

How should we control a network of smart services so that applications requiring low and predictable end-to-end latencies can be hosted on it? In this poster, we will present some models and theory that can be used to control the service capacity and admission control of the smart services so that this goal is achieved. Along with this we will also show how to allow applications to dynamically join and leave the network without sacrificing the predictability of the performance.

Poster 8: Location-aware Workload Prediction for Edge Data Centers
Chanh Nguyen, Umeå University, Sweden

Mobile Edge Cloud (MEC) is a complemented platform to traditional centralized cloud in which IT capabilities are moved closer to user in order to cut down on application-level latency. In MEC, Edge Data Centers (EDCs) are co-located at cellular base stations or wireless access point and provide service to user in its proximity. Due to the bounded coverage area of base station and limited capacity of the EDC in addition to the uncertainty of user location, MEC brings significant challenges in capacity adjustment and planning to optimize cost and resource usage while guaranteeing the expected end-user experience.
To meet these challenges, the resource demand in each EDC is estimated in advance, which is made available for the decision making to efficiently determine various management actions (i.e., allocate/release resource, power on/off server, etc.) to ensure that EDC persistently meets the SLAs of applications hosted, while maximizing resource utilization. To this aim, we are currently working with a Location-aware Workload Prediction method to predict future workload at each EDC in a short time period using a Long Short-Term Memory network. To predict the workload in each EDC, the proposed method not only use its own historical load time series, but also taking into consideration the impact of user mobility by correlating load variations of EDCs in its proximity. The experimental results show that our proposed method are capable of achieving a high accuracy prediction and outperforming the state-of-the-art method.

Poster 9: Cloud Application Predictability through Integrated Load-Balancing and Service Time Control
Tommi Nylander & Marcus Thelander Andrén, Lund University, Sweden

Cloud computing provides the illusion of infinite capacity to application developers. However, data center provisioning is complex and it is still necessary to handle the risk of capacity shortages. To handle capacity shortages, graceful degradation techniques sacrifice user experience for predictability. In all these cases, the decision making policy that determines the degradation interferes with other decisions happening at the infrastructure level, like load-balancing choices. Here, we reconcile the two approaches, developing a load-balancing strategy that also handles capacity shortages and graceful degradation when necessary. The proposal is based on a sound control-theoretical approach. The design of the approach avoids the pitfalls of interfering control decisions. We describe the technique and provide evidence that it allows us to achieve higher performance in terms of emergency management and user experience.

Poster 10: Online Anomaly Detection and Real-Time Causality Recommendation
Mohamad Rezaei, KTH, Sweden

Outlier Detection of machine generated data e.g. logs, and performance data can be a challenging task. First, the rate of generated streams of data can be hard to manage specially when there would be special objectives that need to be handled like real-time objectives. Second, logs usually provide a high dimensional time series data which may not even have bounded dimension in one time series such as database logs. So even if we provide an answer such as the probability of anomaly for a time series, if that time series has too many of features then it might not be possible for the human users e.g. system admins to make sense of why at this point there is an outlier.
In this work, we have created a prototype for anomaly detection system which reads streams of data from a cluster of machines with high number of features that make it hard for humans to explain, and in case of outliers, using conformal prediction in a interactive manner provides an explanation on causality of the detected outlier in that point of time. Using Gradient Boosting Trees with Conformal Prediction as per users input, we sample the data and provide the explanation in form of form of features importance per case of the anomaly model's output. To further improve our model's explanation, we use online domain specific dimensionality reduction on the time series which would be beneficial in cases that number of features are too high or not bounded e.g. database logs. We show that using conformal prediction we can in real-time provide our model explanation independent of the anomaly detection models or data schema.

Poster 11: Server State Estimation Using Event-Based Particle Filtering
Johan Ruuskanen, Lund University, Sweden

Closed-loop control of cloud resources requires there to be measurements readily available from the process in order to utilize the feedback mechanism to form a control law. Sought states might be unfeasible or impossible to measure in real applications; instead they must be estimated. Running the estimators in real time for all measurements however requires a lot of computational overhead. We propose an event-based particle filter approach to capture the internal dynamics of a server with CPU-intensive workload. Preliminary results show some promise as it outperforms analytic estimators derived from stationary equations in service rate estimation for a simulated example. Further we show that for the same example, an event-based sampling strategy outperforms periodic sampling.

Poster 12:Mission Critical Cloud
Per Skarin, Ericsson/Lund University, Sweden

Cloud technology has swiftly transformed the ICT industry and it is continuing to spread. Many ICT applications are suitable for cloud deployment in that they have relaxed timing or performance requirements. In order to take the cloud concepts beyond the ICT domain and apply it to mission critical use cases such as industrial automation, transport and health care we must provide guarantees and predictability. To this end we need new tools and new ways of working. This project attacks this problem from two angles. We will work at developing a cloud infrastructure with a deterministic behaviour, thereby suitable for critical applications. Zero-touch configuration of the cloud based on feedback is a fundamental building block in our approach. Secondly we will showcase the viability of the hardened cloud through mission critical cloud application running in a real data center and operating real-world process, e.g. robotics, unmanned vehicles.

Poster 13: Hybrid Resource Management for HPC and Data Intensive Workloads
Abel Souza, Umeå University, Sweden and Mohamad Rezaei, KTH, Stockholm, Sweden

Traditionally, High Performance Computing (HPC) and Data Intensive workloads have been executed on separate hardware using different tools for resource and application management. With increasing convergence of these paradigms, where modern applications are composed of both types of jobs in complex workflows, this separation becomes a growing overhead. Executing both application classes on the same hardware not only enables hybrid workflows, but can also increase the usage efficiency of the system. While HPC systems are typically managed in a

course grained fashion, allocating a fixed set of resources exclusively to an application, data intensive systems employ a more fine grained regime, enabling dynamic resource allocation and control based on application needs. On the path to full convergence, a useful, less intrusive step is a hybrid resource management system that allows the execution of data intensive applications on top of standard HPC scheduling systems.
In here, we present the architecture of a hybrid system enabling dual-level scheduling for data-intensive and adaptive jobs in HPC infrastructures. Our system takes advantage of resource utilization monitoring to efficiently co-schedule HPC and data intensive applications. The architecture is easily adaptable and extensible to current and new types of distributed workloads, allowing efficient combination of hybrid workloads on HPC resources with increased job throughput with higher overall resource utilization. The architecture is implemented based on the Slurm and Mesos resource managers for HPC and data intensive jobs. Our experimental evaluation in a real cluster based on a set of representative HPC and data intensive applications demonstrate that our hybrid architecture improves resource utilization with 30%, with 35% decrease on queue makespan while still meeting 95% of all deadlines for HPC jobs.

Poster 14: Resource Management in Large Scale Data Centres Using Distributed Optimisation Algorithms
Javad Zarrin, University of Cambridge, UK

Nowadays, Cloud computing is a popular and growing computing paradigm. Clouds are used to offer various applications to the users. A data centre is the back-end infrastructure of the Cloud, enabling to drive Cloud services across public (via the Internet) or private networks. A data centre is composed of clusters of networked computers and storage. Resource management of the data centres is responsible to monitor, control, schedule and allocate resources (e.g., CPUs) to applications. Data centres have significant capital and operational costs stemming from power, cooling and maintenance. Thus, data centres need to operate in a cost-effective manner. However, reports show that for almost any data centre, an average of over 50% of the machines are largely underutilised due to the practice of over-provisioning applications; meaning that resources are allocated to applications based on their most demanding requirements which however are very rare.
Ideally, resource management is expected to perform application scheduling and resource allocation in a way to minimise costs and maximise the performance of applications. The focus of this work is to propose a novel approach to maximise utilisation of data centre resources by applying optimal distributed scheduling with the help of an accurate workload forecasting model. We propose to employ global consensus distributed optimisation algorithms to formulate, design and develop of a multidimensional fully decentralised scheduling mechanism with a scope to provide globally optimal solutions. Furthermore, in order to prevent over-provisioning/under-provisioning issues, we apply machine learning techniques to continuously learn the workload behaviour and accurately predict applications' resource usage. The novelty of this work is on its proposed framework for distributed autonomous resource management based on a unique combination of distributed optimisation and machine learning considering heterogeneity and complexity of large-scale data centres.

Poster 15: Towards Flexible and Adaptive Networks: From Topology Design to Data Center Resource Management
Johannes Zerwas, TU Munich, Germany

Today's communication networks are facing diverse demands which create the necessity for more flexible networks. A broad range of applications along with increased user mobility leads to time-varying requirements for networks. To satisfy those, operators have to adapt networks in timely manner. To accommodate this task, several technologies providing easier adaptation of the network emerged, e.g., Network Virtualization, Software Defined Networking and Cloud Computing. They have the potential to increase the flexibility of the network, i.e., its capability to satisfy new demands, but might also counter-intuitively complicate network design and management by inducing additional configuration possibilities. Moreover, as users expect high quality and seamless services, adaptations must be performed in a timely manner. Thus, highly flexible networks need resource managements with short decision and execution times. With this work we want to promote our research on the optimization of networks for flexibility. In the first step of our research agenda, we start by quantifying network flexibility and try to gain insights on the impact of the network‘s structure, e.g., different data center network topologies, on its provided flexibility. In the second step, while considering these insights, we want to develop algorithms that design and adapt networks to maximize their flexibility. To maintain short adaptation times, we envision to enhance our algorithms based on machine learning and data-driven approaches. Apart from the network's structure, we also want to look deeper into rising challenges when putting flexible network into effect, e.g., we want to focus on the management of physical resources in virtualized environments such as data centers.

Thursday June 14th

8:15 ***Keynote: Deep-Learning-as-a-Service for IoT Systems***
Tarek Abdelzaher, University of Illinois, USA

The Internet of Things (IoT) heralds the emergence of multitudes of computing-enabled networked everyday devices with sensing capabilities in homes, cars, workplaces, and on our persons, leading to ubiquitous smarter environments and smarter cyber-physical "things". The next natural step in this computing evolution is to develop the infrastructure needed for these computational things to collectively learn. Advances in deep learning offer remarkable results but require significant computing resources. This leads to the idea of "Deep-Learning as a Service". The talk will describe back-end services that aim to bring advantages of deep learning to the emerging world of embedded IoT devices. The output of these services are trained neural networks. We discuss several core challenges in training such neural networks for subsequent use in IoT systems. Specifically, while the training process may be resource-intensive, the resulting trained network must be optimized for execution on a resource-limited device while at the same time offering accuracy guarantees. Evaluation results and experiences presented offer encouraging evidence of viability of Deep Learning as a Service for IoT.

9:00 *Interference-aware Resource Management in Cloud Environments*
Anshul Gandhi, Stony Brook University, USA

One of the key performance challenges in cloud computing is the problem of interference, or resource contention, among colocated VMs. While prior work has empirically analyzed interference for specific workloads under specific settings, there is a need for a generic approach to estimate application performance under any interference condition. In this talk, I will present our Markov chain modeling approach to this problem and discuss interesting results under single- and multi-server settings. I will also talk about some of the applications of this work to interference-aware load balancing.

9:20 *Never Late Again! Job-Level Deadline SLOs in YARN*
Subramaniam Venkatraman Krishnan, Microsoft, USA

Modern resource management frameworks for large scale analytics leave unresolved the problematic tension between high cluster utilization and job's performance predictability—respectively coveted by operators and users. We address this in Morpheus [1] a system that: 1) codifies implicit user expectations as explicit Service Level Objectives (SLOs), inferred from historical data, 2) enforces SLOs using novel scheduling techniques that isolate jobs from sharing-induced performance variability, and 3) mitigates inherent performance variance (e.g., due to failures) by means of dynamic reprovisioning of jobs. We validate these ideas against production traces from a 50k node cluster, and show that Morpheus can lower the number of deadline violations, while retaining cluster-utilization, and lowering cluster footprint. We demonstrate the scalability of our implementation by deploying Morpheus on a 2700-node cluster and running it against production-derived workloads. The extensions to the YARN ReservationSystem [2] have been open-sourced as part of Apache Hadoop 2.9.
[1] https://www.usenix.org/conference/osdi16/technical-sessions/presentation/jyothi
[2] https://www.microsoft.com/en-us/research/publication/reservation-based-scheduling-if-youre-late-dont-blame-us-2/

10:10 *Privacy Preservation for WiFi-tracking Data*
Maarten van Steen, University of Twente, The Netherlands

WiFi tracking is becoming increasingly important for domains such as urban planning, crowd management, retailing, and so forth. One of the fundamental problems is that WiFi signals are picked up in public spaces, while at the same time it is virtually impossible to confine detections to only private spaces while still being able to address interesting questions. This distinction is important in light of privacy preservation and the new European General Data Protection Regulation (GDPR). Simply put: WiFi tracking in public spaces such that one adheres to the GDPR is far from trivial. In this talk, I will zoom into our (partly cloud based) approaches to come to WiFi-tracking datasets that are compliant with the GDPR. Next to describing a practical, for the time being most likely sufficient solution, I will also dig into some of the details of the avenues we are exploring to securely come to datasets that fully comply not only with the GDPR, but actually to something that many of us want: knowing what we do without knowing who you are.

10:30 *Synchronization Models for Scalable Machine Learning in the Cloud*
Lixin Gao, University of Massachusetts Amherst, USA
The advances in sensing, storage, and networking technology have created huge collections of high-volume, high-dimensional data. Making sense of these data is critical for companies and organizations to make better business decisions, and brings convenience to our daily life. Recent advances in data mining, machine learning, and applied statistics have led to a flurry of data analytic techniques that typically require an iterative refinement process. However, the massive amount of data involved and potentially numerous iterations required make performing data analytics in a

timely manner challenging. In this talk, we present a series of data parallel frameworks that accelerate iterative machine learning algorithms for massive data.

10:50 *Navigating Heterogeneity and Uncertainty in Private Cloud Scheduling*
Bhuvan Urgaonkar, The Pennsylvania State University, USA

Recent work on fair multi-resource allocation underlies the scheduler of the widely used Apache Mesos cluster manager. We show that heterogeneity (in machines) and dynamism/uncertainty (in machines and workloads) may lead to poor resource utilization in a Mesos cluster. We describe enhancements to the Mesos scheduler that aim to alleviate this problem while preserving fairness properties. Using experiments with synthetic benchmarks and Apache Spark workloads (the latter ongoing), we show that our ideas hold promise.

Discussion 6 (Resarö)
11:10 *Real world challenges in cloud operations*
Narayan Desai, Google, US & Johan Eker, Ericsson Research / Lund University, Sweden

Turning cloud control into an academic discipline while yet addressing existing and relevant problems has proven difficult. Much of the academic work simplifies the problems too much and at the same time the industry fails at posing relevant and coherent questions. In this session we will draw from real-world experience with cloud and DC operation and together with participants discuss how we as a community can formulate relevant topics and back it with data and benchmarking possibilities.

11:10 Discussion 7 (Rindö)
Operational analytics at cloud grade
Eran Raichstein, IBM Research, Israel

In this session we will discuss the challenges, processes and analytic tools available for cloud SREs. We will review available cloud analytic techniques and discuss the challenge of "closing the loop". Efficient data acquisition @scale; persistency and data-curation; analytics and insights; adjustment application; and back to acquisition. We will focus on multi-layer network analytics as a potential impact path, and discuss the place of manual vs. automated processes in such a challenge.

11:10 Discussion 8 (Storskär)
Making big data processing just work
Bogdan Ghit, Databricks, The Netherlands & Cristian Klein, Umeå University, Sweden

Modern data analytics tools for large scale data processing have paved the way towards understanding the world around us through personalized medicine, machine learning, and genomics. However, despite many years of work both from industry and academia, big data tools are still rather difficult to employ in order to setup a reliable, production pipeline (compare with web apps for instance). The promise of automatic deployment in the cloud with out of the box performance guarantees has not been fully achieved as we are still required to tune a wide parameter space, while integrating a large number of subsystems both vertically and horizontally (the data processing stack from the high-level API down to the storage layer). The main discussion points are as follows:
1. Which are the current limitations of the big data tools?
2. What can we do better to improve them?
3. How can the cloud community help?.

11:10 Discussion 9 (Huvudskär)
A Tale of Distribution: Continuous and Event-based Software Engineering Strategies for Cloud Management
Martina Maggio, Lund University, Sweden

Cloud computing infrastructures comprise multiple layers and many different components. Load balancers receive incoming requests. Servers handle the requests generating responses that should be forwarded to clients. Autoscalers handle the online addition and removal of servers to adapt to the current load. Monitors keep track of the current status of the components and react in case of failures.
Managing the cloud infrastructure requires feedback from the different components and some logic to handle multiple events. The distributed nature of the cloud infastructure requires the use of event-triggered responses to handle changes at runtime, e.g., component failures. This triggers the need for software engineering solutions that are capable of integrating continuous monitoring and management, and event-based responses. The discussion session covers some of the challenges in this complex landscape as well as hints to potential solutions for these challenges.

11:10 Discussion 10 (Svavelsö)
Distributed intelligence: How can Artificial Intelligence help with the challenges of managing and controlling dynamic Fog/Edge compute platforms? Joe Butler & Thijs Metsch, Intel Labs, Ireland

In this session, we want to discuss how topics from the field of Artificial Intelligence (Note: not exclusively focusing on Deep Learning and Neural Networks btw.) help with today's challenges of managing and controlling Fog/Edge compute platforms. Challenges that we want to address, include - but are not limited to – dynamicity of systems in which entities (e.g. cars, trains, etc.) are moving, the workloads that require Edge compute platforms, distribution of resources to meet latency requirement, and heterogeneity of the resource to match various workloads. All these challenges make the general topic of mapping workloads requirements to the available resource in a dynamic environment more challenging and provide scope to look back (and into the future) on what the field of Artificial Intelligence has to offer to solve them.

13:45 **Keynote*: Building the Warehouse Scale Computer***
John Wilkes, Google, USA

Imagine some product team inside Google wants 100,000 CPU cores + RAM + flash + accelerators + disk in a couple of months. We need to decide where to put them, when; whether to deploy new machines, or re-purpose/reconfigure old ones; ensure we have enough power, cooling, networking, physical racks, data centers and (over longer a time-frame) wind power; cope with variances in delivery times from supply logistics hiccups; do multi-year cost-optimal placement+decisions in the face of literally thousands of different machine configurations; keep track of parts; schedule repairs, upgrades, and installations; and generally make all this happen behind the scenes at minimum cost.

And then after breakfast, we get to dynamically allocate resources (on the small-minutes timescale) to the product groups that need them most urgently, accurately reflecting the cost (opex/capex) of all the machines and infrastructure we just deployed, and monitoring and controlling the datacenter power and cooling systems to achieve minimum overheads - even as we replace all of these on the fly.
This talk will highlight some of the exciting problems we're working on inside Google to ensure we can supply the needs of an organization that is experiencing (literally) exponential growth in computing capacity.

14.30 Discussion 11 (Resarö)
Cloud Computing for an AI First future
Geoffrey Fox, Indiana University Bloomington, USA

Artificial Intelligence is a dominant disruptive technology affecting all our activities including business, education, research, and society. Further, several companies have proposed AI first strategies. The AI disruption is typically associated with big data coming from edge, repositories or sophisticated scientific instruments such as telescopes, light sources and gene sequencers. AI First requires mammoth computing resources such as clouds, supercomputers, hyperscale systems and their distributed integration. This discussion will examine the driving applications and their implication for hardware and software infrastructure.

14.30 Discussion 12 (Rindö)
Hybrid cloud orchestration, can we make it simple?
Gal Hammer, Red Hat, Israel & Luis Tomas, Red Hat, Spain

Nowadays Hybrid (multi) clouds are becoming a reality, however managing them as well as fully take advantage of them is still quite complex. Hybrid cloud are compose of too many different layers and building blocks. Not only connecting different clouds (e.g., private OpenStack based clouds with public AWS or GCE resources) but also different resources within the same provider: baremetal servers, Virtual Machines, Containers, networking equipment (both physical and virtual). Each one with its own management system. And on top of that there is usually an orchestrator (or several) to deploy/control infrastructure and applications' lifecycle. Consequently, there is a clear need for automation and abstraction models that would enable a simple to use and manage system.

14.30 Discussion 13 (Storskär)
Software Defined Infrastructure - Opportunities and challenges
Mats Jonsson, Saab AB, Sweden

Saab is in the middle of designing and building an almost entirely software defined datacenter, scalable from small distributed installations to warehouse size. Designed to be continuously on-line, fault tolerant, gracefully degrading, highly performant, resilient, compliant to strict security requirements, upgradeable and updateable, we've had to resolve a wide range of conflicts and issues. We will share some entertaining misstakes and failures made along the way, and some of the more interesting and unexpected successes.
We will tell the tale of how we

- got involved in the standardization of programmable dataplanes,
- almost bricked a datacenter with Chaos Engineering

- Lost control of distributed storage and don't know to put the pieces back together
- Use and plan to expand end-to-end in-band telemetry in our SDN to drive Lineage Driven Fault Injection (LDFI) and large scale Chaos Engineering to the next level
- plan to take advantage of short startup times for unikernels/microservices/FaaS and do datacenter scale garbage collect and load balancing

and why Linus Torvalds was correct when he said "f*ck you nVidia!"

14.30 Discussion 14 (Huvudskär)
Edge Clouds: Limitations and Applications
Ahmed Ali-Eldin, Umeå University, Sweden / UMass, Amherst, US & Farah Ait Salaht, INRIA-Grenoble, France

While edge clouds have garnered recently the interest of the research community, recent results show that there are limitations with today's technologies on what can be achieved showing that in some common cases the latency of the first-aggregation stage dominates the end-to-end latency [1]. In addition, there is an increased cost factor for deploying such an infrastructure that not just entails the extra equipment, but also the logistics of deploying distributed resources that require maintenance and fault tolerance at a new scale. While some hope to get down the delay to just the speed of light using 5G [2], this is unlikely in the near future and it is hard to distill the science from the marketing for some of these claims [3,4]. In this session, we will discuss the current State-of-the-Art, what applications will justify the tremendous cost of building an edge cloud, and what problem does the research community need to focus on right now.

[1] Hadžić, Ilija, Yoshihisa Abe, and Hans C. Woithe. "Edge computing in the ePC: a reality check." In Proceedings of the Second ACM/IEEE Symposium on Edge Computing, p. 13. ACM, 2017.
[2] http://money.cnn.com/2018/02/05/technology/business/5g-internet-of-skills/index.html
[3] https://medium.com/@miccowang/5g-when-will-we-see-it-7c436a4ad86c
[4] https://www.computerworld.com.au/article/641759/don-t-believe-hype-we-re-long-way-from-5g/

14.30 Discussion 15 (Svavelsö)
Practices and challenges for building cloud-native applications with FaaS,
Philipp Leitner & Joel Scheuner, Chalmers / University of Gothenburg, Sweden

"Serverless" is a major trend in cloud computing. Serverless cloud applications are deployed to infrastructure components that are entirely transparent to the application developer. One core implementation of this model are Function-as-a- Service (FaaS) offerings, as most prominently exemplified in AWS Lambda. Unfortunately, building non-trivial cloud-native applications on top of FaaS systems is a difficult proposition: existing tooling is in its early stages, there is limited support for integration testing and debugging, and the compositional programming model underlying serverless applications is different enough to standard IaaS or PaaS that many developers struggle. In last year's Cloud Control workshop, one participant characterised AWS Lambda and similar services to "Assembly for Cloud-Native, in desperate need for higher-level programming constructs". In this session, we want to re-visit and continue last year's discussion on practices and challenges for actually building applications using FaaS.

Friday June 15th

8.15 ***Keynote**: Shared Memory and Disaggregated Memory in Virtualized Clouds*

Ling Liu, Georgia Institute of Technology, USA

Cloud applications are typically deployed using the application deployment models, comprised of virtual machines (VMs) and/or containers. These applications enjoy high throughput and low latency if they are served entirely from main memory. However, when these applications cannot fit their working sets in real memory of their VMs or containers, they suffer severe performance loss due to excess memory paging and thrashing. Even when unused host memory or unused remote memory are present in other VMs on the same host or across the cluster, these applications are unable to benefit from those idle host/remote memory. In this keynote, I will first revisit and examine the problem of memory imbalance and temporal usage variations in virtualized Clouds and discuss the potential benefits of dynamically managing and sharing unused host memory and unused remote memory. Then I will describe some systems solutions for exploiting shared memory and disaggregated memory transparently, opportunistically, and non-intrusively. and present our initial results for efficient sharing of host and remote memory. The talk will conclude with a discussion on integrating shared memory and disaggregated memory management as an integral part of Cloud control for hosting big data and machine learning workloads in Cloud datacenters.

9:00 *Harvesting Randomness to Optimize Distributed Systems*

Siddharta Sen, Microsoft, USA

We present a vision for optimizing cloud systems without disrupting them. The idea is to leverage the natural randomness and information inherent in these systems, and the logs they already emit, to ask counterfactual questions like: "What would happen if I changed my system in this way?". We show how techniques from reinforcement learning can be combined with ideas from systems to answer such questions, while minimally affecting the live system. Our methodology is general: it does not rely on domain knowledge or accurate modeling. We apply our methodology to three cloud infrastructure systems in Azure: a cluster controller monitoring machine health, a geo-distributed proxy routing user requests to cloud services, and a replica selector load balancing requests to a storage layer. In each case, we show that we can counterfactually evaluate arbitrary policies, without ever deploying them, and obtain estimates that closely match the ground truth. We use this methodology to synthesize new (machine-learned) policies that improve significantly upon the production system baselines.

9.20 *Networking as a First-Class Cloud Resource*

Rodrigo Fonseca, Brown University, USA

Tenants in a cloud can specify, and are generally charged by, resources such as CPU, storage, and memory. There are dozens of different bundles of these resources tenants can choose from, and many different princing schemes, including spot markets for left over resource. This is not the case for networking, however. Most of the time, networking is treated as basic infrastructure, and tenants, apart from connectivity, have very little to choose from in terms of network properties such as priorities, bandwidth, or deadlines for flows.

In this talk I look into why that is, and whether networking could be treated as a first-class resource. The networking community has developed plenty of mechanisms for different networking properties, and programmable network elements enable much more fine-grained control and allocation of network resources. We argue that there may be a catch-22, as tenants can't specify what they want, and providers, not seeing different needs, don't provide different services, or charge differently for these services. I will discuss a prototype we have designed with the Massachusetts Open Cloud project, which provides a much more expressive interface between tenants and the cloud for networking resources, improving efficiency, fostering innovation, and even allowing for a marketplace for networking resources.

10.10 *WIP: Exploiting Heterogeneity to Minimize Tail Latency*

Timothy Zhu, The Pennsylvania State University, USA

Tail latency (e.g., 99th percentile) is an important metric when operating user-facing services. Simultaneously, cost predictability is desirable for many companies using the cloud. In our ongoing work, we are investigating techniques for minimizing the tail latency of user-facing services under fixed hourly costs. Specifically, we are studying the potential of using a heterogeneous collection of resources (e.g., fast expensive servers mixed with cheap slow servers) to trade-off speed and cost. While naively using heterogeneous collections of servers can balance cost and average latency, we find that extra care must be taken to address tail latency, which is often dominated by the performance of the slow servers. By carefully scheduling requests, our preliminary results suggest that utilizing heterogeneity can nevertheless bridge the performance gap between different server types, even for tail latency.

10:30 *Decentralized Self-Adaptation for Elastic Data Stream Processing*

Valeria Cardellini, University of Rome Tor Vergata, USA

Data Stream Processing (DSP) applications should be capable to efficiently process high-velocity continuous data streams by elastically scaling their execution on multiple geo-distributed computing nodes, so to deal with highly variable workloads. In this talk, we present a hierarchical distributed architecture for the autonomous control of elastic DSP applications. It revolves around a hierarchical approach, with two control layers that work at different granularity and time scale. We present distributed self-adaptation policies for the elasticity of DSP operators, including a popular threshold-based approach and two reinforcement learning solutions. We also consider the multi-level elasticity of DSP applications, by exploiting infrastructure-level elasticity in addition to application-level elasticity. Thanks to its decentralized adaptation features, the proposed solution can efficiently run in geo-distributed Cloud and Fog systems.

10:50 *Back to the Future: Towards Timely Bounded Cloud Communication*
Patrick Eugster, University of Lugano, Switzerland

While datacenter networks keep getting faster, they typically still offer best effort guarantees for communication. While this suffices for many applications or major parts thereof, coordination and synchronization tasks are still unnecessarily complex due to the absence of bounds on communication delays. In this talk I will present first results towards fulfilling a long envisioned concept: timely bounded communication in commodity networks.

11:10 Discussion 16 (Resarö)
Cloud versus in-network-processing for latency-critical industrial control operations
Anders Robertsson, Lund University, Sweden & Klaus Wehrle, RWTH Aachen University, Germany

Abstract delayed due to late merger of discussion session topics.

11:10 Discussion 17 (Rindö)
Customised Data as a Service for multiple and conflicting data intensive applications in cloud/edge environments
Monica Vitali, Politecnico di Milano, Italy

IoT sensors are spread all over the world collecting data that might be useful to several customers for different purposes. The Data as a Service paradigm can be used to provide access to data to the different data consumers. In this way, the user can access data without worrying about the data management. The DaaS, instead, manages the data taking decision on copying and moving data from one location to another, either in the edge of the network and in the cloud. In doing so, the DaaS might consider different requirements of the users in terms of Quality of Service. In order to achieve this goal, the DaaS needs to be customised for each user, which can express its requirements through an SLA. The goal of this session is to investigate challenges and opportunities of managing a customised DaaS where the requirements of the different users and applications accessing the data source might be in contrast with each other.

11:10 Discussion 18 (Storskär)
Managing Container Clouds: Challenges and Opportunities
Ali Kanso, IBM, US

Linux containers are shaping the way we are executing workloads in the cloud . Every major cloud provider is offering a managed container infrastructure where the cloud tenants can push and execute their workload. Moreover, containers are enablers for other compute models such as Serverless and event driven computing. In this session we will shed some light on the challenges of offering containers in a public cloud. We will mainly focus in the issues of clustering containers, Securing containers, fast scaling and auto-scaling, failure recovery and preventing cascading failures. We will also touch upon the future of container technologies and the direction to which the industry is heading as well as the areas where research contributions are highly needed.

11:10 Discussion 19 (Huvudskär)
Data plane programmability for flexible cloud networking
Andreas Kassler, Karlstad University, Sweden

The performance and functionality of the datacenter network is an important aspect for the overall performance and scalability of Cloud based applications and services. This is because the datacenter network is at the heart of the cloud infrastructure and must provide a resilient, high capacity and low latency forwarding fabric. As of now, most datacenters use commodity off-the-shelf servers and switches and run standard operating systems. However, traditional networking equipment and protocols have been optimized for wide area networks with the goal to maximize throughput and latencies are several orders of magnitudes higher than inside a datacenter.
For next generation datacenter networks, increasing their functionality and flexibility is as important as controlling and minimizing the datacenter network latency. Especially the recent emergence of programmable switch and network interface card architectures together with compiler support for the P4 programming language have led to many great research ideas. Examples include novel in-network load-balancing (e.g. Hula), in-band network telemetry for real-time data plane monitoring (INT), in-network caching for key-value stores (NetCache) or stateful server load-balancing in the data plane (SilkRoad). The goal of this session is to discuss recent advances in programmable data plane architectures for next generation datacenter networks that facilitate a much tighter integration of compute and networking.

11:10 Discussion 20 (Svavelsö)
Single System Image and distributed OS: Has the time come?
Daniel Turull & Mina Sedaghat, Ericsson Research, Sweden

Can we revisit the ideas of earlier Single System Image (SSI) and distributed operating system work and adapt them for the cloud, i.e. to make datacenters easier to deploy and manage? Can we get rid of thinking about "boxes" as execution abstractions (VMs, containers) and instead focus on an SSI abstraction? Could we convert the networking within a DC to an IPC space? How can the cloud scale from a single node to millions of nodes without of the need of human interaction? This session will focus in different advantages and disadvantages of distributed control mechanisms and how modern clouds could benefit from them.

13.45 *Reducing the Cost of Exploring Deployment Configurations in the Cloud*
Luis Rodrigues, INESC-ID, IST, Ulisboa, Portugal

Finding the best deployment for applications in the cloud is extremely hard due to the diversity of options available that result in large search spaces: cloud providers offer several different machines types and sizes and complex application have many configuration parameters. One would like to build models that help in selecting the right deployment options in an automated manner. However, such models cannot be built without experimenting multiple configurations, and building the model can be a slow and costly task. This talk reports on current research directions that aim at reducing the exploration cost, by selecting carefully which configurations to experiment. The selection process combines the expected contribution of the experiment to improve the accuracy of the model and the actual cost of running each experiment.

14:05 *Massivizing Computer Systems: a Vision to Understand, Design, and Engineer Computer Ecosystems through and beyond Modern Distributed Systems*
Alexandru Iosup, Vrije Universiteit Amsterdam, The Netherlands

Our society is digital: industry, science, governance, and individuals depend, often transparently, on the interoperation of large numbers of distributed computer systems. Although the society takes them almost for granted, these computer ecosystems are not available for all, may not be affordable for long, and raise numerous other research challenges. Inspired by these challenges and by our experience with distributed computer systems, we envision Massivizing Computer Systems, a domain of computer science focusing on understanding, controlling, and evolving successfully such ecosystems. Beyond establishing and growing a body of knowledge about computer ecosystems and their constituent systems, the community in this domain should also aim to educate many about design and engineering for this domain, and all people about its principles. This is a call to the entire community: there is much to discover and achieve.

14:25 *Experimental Evaluation of Amazon EC2 Spot Instances*
Thomas Fahringer, University of Innsbruck, Austria

Amazon EC2's spot instances (SIs) represent a competitive Cloud resource in terms of price compared to reliable and fixed price options. The drawback, however, is that SIs may not always be available and they can be revoked at any given time. In this presentation, we describe a comprehensive experimental evaluation for EC2 SIs to characterize their performance and behavior in three different regions each of which in a different continent. We describe the life cycle of SIs with the most important phases of an SI, introduce the most relevant events that can prevent a user from obtaining SIs, and draw important conclusions that can be exploited by the research community to effectively use the spot market. Our results reveal the fulfillment rate of requests for SIs, waiting time until requested SIs become fulfilled, details about the interruption rate of SIs, and how long SIs run before being interrupted.

14:45 *Closing*
Erik Elmroth, Umeå University, Sweden

We will summarize the workshop and look a bit ahead towards future events. The 15th Cloud Control Workshop is planned to take place at Umeå University, Sweden, some days during June 16-20, 2019. It will be co-located the 16th IEEE International Conference on Autonomic Computing (ICAC 2019) (which is now merged with the IEEE International Conference on Cloud and Autonomous Computing, ICCAC) and The 13th IEEE International Conference on Self-Adaptive and Self-Organizing Systems (SASO 2019). Although co-located with regular "paper oriented" conferences, we will for the Cloud Control Workshop strive to maintain the form and spirit established during the past few years. One obvious difference is that all participants will not be locked-in on the same remote hotel. Instead, we will strive to turn the whole of Umeå into a vibrant Cloud Control meeting place for a few days before midsummer 2019.