Program

7th Cloud Control Workshop

Nässlingen, Stockholm Archipelago, Sweden

June 9 – 11, 2015

Tuesday June 9

8.20 Gathering at bus outside Arlanda, Terminal 4

8.30 Bus departure

10.10 Boat from Åsättra at Ljusterö to Nässlingen

10.30 Coffee in lounge outside conference room

11.00 Workshop Introduction and introductions from participating organizations
Erik Elmroth, Umeå University

12.00 Lunch

Session Chair: Maria Kihl, Lund University

13.00 **Keynote: *Hybrid Cloud Control: Scheduling workloads across enterprise datacenters and the off-premise cloud***
Anne Holler, VMware, Palo Alto

14.00 *Fairness in Data Stream Processing under Overload*
Eva Kalyvianaki, City University London

14.20 Discussion 1: *Ultimate Test Facility for Cloud Control Research*
Tor Björn Minde, Ericsson Research, Kista & Erik Elmroth, Umeå University
Notes: Ahmed Ali-ElDin, Umeå University

15.30 Coffee

Session Chair: Olov Schelén, Luleå University of Technology

16.00 **Keynote: *High-resolution representation of the infrastructure and service landscapes, why and how***
Joe Butler, Intel, Leixlip

16.40

Discussion 2:	Discussion 3:	Discussion 4:
Managing containers running in the cloud Anne Holler, VMware, Palo Alto Notes: Mina Sedaghat, Umeå University	*Opportunities and Challenges for Distributed Cloud and Telecom* Johan Eker, Ericsson Research, Lund & James Kempf, Ericsson Research, San Jose / TU Berlin Notes: Amardeep Mehta, Umeå University/William Tärneberg, Lund University	*Cloud Control for Realistic Clouds* Karl-Erik Årzén, Lund University Notes: Johan Tordsson, Umeå University

18.00 End of session

19.00 Dinner

Wednesday June 10

Session Chair: Anders Robertsson, Lund University

8.20 **Keynote: *A Control-Theoretic Approach for Dynamic Adaptive Video Streaming over HTTP***
Bruno Sinopoli, Carnegie Mellon University, Pittsburgh

9.20 *Model-Based Deadtime Compensation of Virtual Machine Startup Times*
Mannfred Dellkrantz, Lund University

9.35 *Dynamic Resource Allocation for Service Response Time Control*
Ewnetu Bayuh Lakew, Umeå University

9.50 Coffee

Session Chair: Per-Olov Östberg, Umeå University

10.20 *Mobile Cloud Networking; beyond NFV*
Andy Edmonds, ZHAW, Zürich

10.40

Discussion 5:	Discussion 6:	Discussion 7:
Cloud Control in a new Era: Opportunities and challenges that stem from disaggregated hardware resources Azimeh Sefidcon & Fetahi Wuhib, Ericsson Research, Kista Notes: Jakub Krzywda, Umeå Univ.	*Replacing the human data center operator with a script!* Thijs Metsch & Joe Butler, Intel, Leixlip Notes: Olumuyiwa Ibidunmoye, Umeå University	*Dealing with uncertainty via stochastic optimization and randomized algorithms* Alessandro Papadopoulos and Martina Maggio, Lund University Notes: Jonas Dürango, Lund University

12.00 Lunch

Session Chair: Rolf Stadler, KTH, Stockholm

13.00 **Keynote: *Large-scale cluster management at Google with Borg***
John Wilkes, Google, Mountain View

14.00 Discussion 8 (Panel): *The Future Challenges for Cloud Control*
Thomas Bohnert, ZHAW, Zürich
Joe Butler, Intel, Leixlip
Anne Holler, VMware, Palo Alto
Azimeh Sefidcon, Ericsson Research, Kista
Bruno Sinopoli, Carnegie Mellon University, Pittsburgh
John Wilkes, Google, Mountain View

Notes: Ahmed Ali-ElDin, Umeå University

15.30 Coffee

16.00 Social outdoor activities

19.00 Dinner

Thursday June 11

Session Chair: Christer Åhlund, Luleå University of Technology

8.30 *Zero-touch cloud service assurance enabled by analytics*
Andreas Johnsson, Ericsson Research, Kista

8.50 *Bringing Probablistic Guarantees to Autoscaling Algorithms*
Ahmed Ali-ElDin, Umeå University

9.05 *Performance Anomalies in Systems at Scale: Detection, Diagnosis, and Research Trends*
Olumuyiwa Ibidunmoye, Umeå University

9.20 *On Control-Theoretical Approaches to Cloud Spot Pricing*
Zheng Li, Lund University

9.35 *Dynamic Application Placement in the Telco-cloud*
Amardeep Mehta, Umeå University

9.50 Coffee

Session Chair: Tor Skeie, Simula Research, Fornebu

10.20 *MITOSIS: distributed autonoMIc managemenT Of ServIce compoSitions*
Giovanni Toffetti, ZHAW, Zürich

10.40	Discussion 9:	Discussion 10:	Discussion 11:
	Analytics for Cloud Control and Management Rolf Stadler, KTH, Stockholm & Andreas Johnsson, Ericsson Research, Kista Notes: Olumuyiwa Ibidunmoye, Umeå University	*Self-* applications versus cloud-vendor management services, are they worth the effort?* Andy Edmonds & Giovanni Toffetti, ZHAW, Zürich Notes: Luis Tomás, Umeå University	*Self-adaptive network architecture for InfiniBand based HPC clouds* Feroz Zahid & Ernst Gunnar Gran, Simula Research, Fornebu Notes: Per-Olov Östberg, Umeå University

12.00 Lunch

13.00 *The Wallenberg Autonomous Systems Program (WASP)*
Karl-Erik Årzén, Lund University

13.15 *Concluding discussions*
Erik Elmroth, Umeå University

Loose ends?
Feedback and follow-up on this workshop
The Future of Cloud Control Workshops

14.00 End of session – Walk to harbour

14.20 Boat departure from Nässlingen

14.40 Bus departure from Åsättra at Ljusterö

16.30 Latest expected arrival time at Arlanda airport, Terminal 4

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

Tuesday, June 9

11.00 – 12.00 Workshop Introduction and introductions from participating organizations
Erik Elmroth, Umeå University

13.00 – 14.00 Keynote: ***Hybrid Cloud Control: Scheduling workloads across enterprise datacenters and the off-premise cloud***
Anne Holler, VMware, Palo Alto

Private datacenters are typically provisioned to accommodate projected peak resource demand. When peak is significantly higher than average demand, such datacenters are inefficiently used, and when peak demand increases every year, continuing to provision enterprise datacenters for peak is difficult to sustain. The hybrid cloud model, in which off-premise resources are purchased to extend enterprise datacenters, promises to reduce the continued growth in capital expenses needed to fully accommodate peak. To exploit the hybrid cloud model, customers would like to use their sunk-cost enterprise datacenter resources when possible and to only use cloud resources when their enterprise datacenters are too busy and only for workloads complying with their company's hybrid cloud usage policies. Customers would also like workload placement decisions to be automated, to allow seamless operation at scale. This talk presents HybridDRS, a resource scheduler that performs seamless burst scheduling from customers' enterprise datacenters to an off-premise cloud. HybridDRS makes placement decisions that balance placement constraints, resource availability, and cost across the candidate placement targets.

14.00 – 14.20 Fairness in Data Stream Processing under Overload
Eva Kalyvianaki, City University London

Federated stream processing systems (FSPSs), which utilise nodes from multiple independent domains, can be found increasingly in multi-provider clouds. To pool resources from several sites submitted queries are executed collaboratively by different sites. When supporting many users, queries may exhaust available processing resources, thus requiring constant load shedding. Given that individual sites have autonomy, it is an open challenge how to ensure global fairness on queries' processing quality. We describe THEMIS, a FSPS for resource-starved, multi-site deployments. It executes queries in a globally fair fashion and provides users with constant feedback on the experienced processing quality for their queries.

14.20 – 15.30 Discussion 1: *Ultimate Test Facility for Cloud Control Research*
Tor Björn Minde, Ericsson Research, Kista & Erik Elmroth, Umeå University

SICS Swedish ICT is planning to establish a large-scale cloud and datacenter research facility. The facility will support research spanning from building structure, energy, and cooling supply to cloud and analytics software platforms, including cloud management systems. After a short introduction to the current plans, the discussion will focus on what such a large-scale test facility ideally should look like to support research in cloud control, to answer questions like what would we like to measure, what should be possible to configure, what level of control of the facility do we need, what workloads would we like to see on the systems, etc. The discussions will have a starting point in what research we would like to perform and, if sky is the limit, what experiments we ideally would like to perform.

16.00 – 16.40 Keynote: High-resolution representation of the infrastructure and service landscapes, why and how
Joe Butler, Intel, Leixlip

Workload placement, capacity planning and run-time adjustment typically operate with reference to an asset-list style resource inventory. While this approach has proven effective in dealing with bounded service and workload ranges, it falls short when more complex service and infrastructure combinations become relevant. The trends toward micro-service architectures and heterogeneous infrastructures including storage, compute accelerators, and specialisations expose the need for a richer view of the landscape. Key issues such as efficiency and return-on-investment are more readily addressed when the landscape representation incorporates contextual details such as capacity, utilization over time and effective performance. Intel Labs are exploring approaches to automatically deriving and maintaining detailed landscape views and using them as a basis for analytics-supported intelligent orchestration.

16.40 – 18.00 Discussion 2: Managing containers running in the cloud
Anne Holler, VMware, Palo Alto

Enterprise datacenters have become predominantly virtualized, and cloud infrastructure as a service is typically provided using virtual machine encapsulation. Recently, interest has grown in creating or rearchitecting enterprise applications as collections of microservices running in containers, for easy deployment and upgrade. In this session, we'll discuss the implications of managing applications running in containers in the cloud, including how such applications are structured, how resources are allocated to them, and how they are managed on an on-going basis.

16.40 – 18.00 Discussion 3: Opportunities and Challenges for Distributed Cloud and Telecom

Johan Eker, Ericsson Research, Lund & James Kempf, Ericsson Research, San Jose / TU Berlin

The research work on 5G just getting underway provides a natural starting point for introducing virtualization and cloud computing management techniques into mobile networks. In wired networks, the transition to virtualization and cloud is already underway through the European Telecommunications Standards Institute (ETSI) Network Functions Virtualization (NFV) initiative. A natural evolution for operators is to re-architect their central offices and base station sites into mini data centers, leading to a highly distributed cloud architecture. A distributed cloud architecture, in which numerous smaller localized data centers predominate while a few mega data centers are still available for long-term storage and latency-insensitive applications, seems the best match for supporting telecommunications infrastructure as a service (TIaaS). Two high level drivers for the cloudification of the mobile network (core + RAN) are: improving the latency for sensitive applications and enabling applications with large locally relevant data sets. Moreover, improved operator manageability of the Evolved Packet Core (EPC) provided through virtualization of the EPC network functions highlights how additional gains may be achieved. Finally, for service development and deployment agility, the Network Management System and Operations Support System (OSS) of the future seems likely to be cloud-based, and to feature open interfaces so developers from a wide variety of third party providers can program new services. Such interfaces have been extremely successful in fostering a broad and growing offering of services on the Internet, and the hope is that they will achieve the same effect for telecommunications services. In this discussion section, we will discuss the promise and challenges of creating TIaaS.

16.40 – 18.00 Discussion 4: Cloud Control for Realistic Clouds

Karl-Erik Årzen, Lund University & Johan Tordsson, Umeå University

Feedback-based "cloud control" has mostly been proposed for web service applications, with feed-based elasticity/auto-scaling as the most common example. However, realistic clouds/data centers contains a mixture of batch jobs (Hadoop, MPI, ...) and interactive applications (VM encapsulated servers, business logic, storage, caches,...). For these systems resource managers such as Mesos and Omega have been developed. The topic of this discussion session is what the possibilities are of introducing feedback control in these types of systems. What are the control objective, sensors, and actuators?

Wednesday, June 10

8.20 – 9.20 Keynote: A Control-Theoretic Approach for Dynamic Adaptive Video Streaming over HTTP

Bruno Sinopoli, Carnegie Mellon University, Pittsburgh

User-perceived quality-of-experience (QoE) is critical in Internet video applications as it impacts revenues for content providers and delivery systems. However, there is little support in the network for optimizing such measures and bottlenecks could occur anywhere in the delivery system (e.g., content delivery network, user home, ISP). Consequently, a robust bitrate adaptation algorithm in client-side video players is critical to ensure good user experience. This is a key problem for the Internet video ecosystem as new standards for dynamic adaptive streaming over HTTP (DASH) are emerging.

Previous studies by leading commercial providers (e.g., YouTube, NetFlix, Conviva) and academic efforts have shown key limitations of state-of-art commercial solutions and proposed a broad spectrum of heuristic fixes. Despite the emergence of several proposals, there is still a distinct lack of consensus on (a) how best to design this client-side bitrate adaptation logic (e.g., use rate estimates vs. buffer occupancy); (b) how well specific classes of ap- proaches will perform under diverse operating regimes (e.g., high variability); or (c) how they actually balance different Quality-of-Experience (QoE) objectives (e.g., startup delay vs. buffering).

In this work, we formalize the DASH problem through the "lens" of control theory and develop a principled control-theoretic model to reason about a broad spectrum of adaptation strategies. We propose a novel model predictive control algorithm that can optimally combine bandwidth and buffer-size information to significantly outperform traditional approaches in common conditions. Finally, we describe a practical implementation of this algorithm in the industry-reference video player DASH.js to confirm the validity of the approach in real world scenarios.

9.20 – 9.35 Model-Based Deadtime Compensation of Virtual Machine Startup Times

Mannfred DellKrantz, Lund University

Scaling the amount of resources allocated to an application according to the actual load is a challenging problem in cloud computing. The emergence of autoscaling techniques allows for autonomous decisions to be taken when to acquire or re- lease resources. The actuation of these decisions is however affected by time delays. Therefore, it becomes critical for the autoscaler to account for this phenomenon, in order to avoid over- or under-provisioning. This paper presents a delay-compensator inspired by the Smith predictor. The compensator allows one to close a simple feedback loop around a cloud application with a large, time-varying delay, preserving the stability of the controlled system. It also makes it possible for the closed-loop sys- tem to converge to a steady-state, even in presence of re- source quantization. The presented approach is compared to a threshold-based controller with a cooldown period, that is typically adopted in industrial applications.

9.35 – 9.50 Dynamic Resource Allocation for Service Response Time Control

Ewnetu Bayuh Lakew, Umeå University

Services hosted in the cloud have become indispensable for many business and personal processes, making performance of these services a key issue. However, today, cloud infrastructure providers either do not offer any performance guarantees or prefer coarse-grained and static Virtual Machine (VM) provisioning with fixed sizes. This usually leads to either under- or over-provisioning of resources and consequently result in service level objective violations or inefficient resource utilization. In order to guarantee strict performance requirements of services and increase resource utilization, resources should be allocated dynamically using fine-grained resource allocation techniques as required by services. This talk briefly presents our recent research on controlling service response time targets by dynamically adjusting the right amount of resources required by each service.

10.20 – 10.40 Mobile Cloud Networking; beyond NFV

Andy Edmonds, ZHAW, Zürich

The virtualisation of telcom workloads is best represented by ETSI NFV efforts. In this presentation, we present Mobile Cloud Networking that pre-dates the work on NFV yet shares a common vision and compliant architecture. Details of composed multi-tenant services will be shown and a deep dive on how the vision is technically realised. Aspects of full-lifecycle multi-domain service composition, orchestration and interoperability will form part of the talk.

10.40 – 12.00 Discussion 5: Cloud Control in a new Era: Opportunities and challenges that stem from disaggregated hardware resources

Azimeh Sefidcon & Fetahi Wuhib, Ericsson Research, Kista

Traditional server architecture boundaries have been broken and hardware resources are becoming independent units within a rack. This trend is widely known as hardware disaggregation. Beyond the obvious impact such a trend has on the hardware layer of a cloud stack, and in order to take the most out of such flexible hardware architecture, a proper integration with the remaining cloud stack must be in place. This session will try to identify and discuss opportunities and key research challenges this trend presents to the cloud research community.

10.40 – 12.00 Discussion 6: Replacing the human data center operator with a script!

Thijs Metsch & Joe Butler, Intel, Leixlip

Feedback control and Complex Event Processing (CEP) systems have not taken over the world from human data center operators for controlling clouds. Specialized systems exist (policy driven) and have proven their value, but the question arises: why can't they be lifted out of their particular domains? Don't they scale? Are inputs missing? Are we still trying to find the right (data science) algorithms? During this discussion session we want to foster the hive-mind of people in the room and tackle these topics and evaluate why this is the case or not. And if it is the case take a close look at shortcomings of existing systems to lead to road to the future solutions.

10.40 – 12.00 Discussion 7: Dealing with uncertainty via stochastic optimization and randomized algorithms
Alessandro Papadopoulos and Martina Maggio, Lund University

Uncertainty is a ubiquitous problem in all the fields of engineering. Every mathematical model is just an approximation of the reality, and, as such, it carries also an inherent uncertainty. In addition, in more complex applications, such as cloud computing, there is a number of unpredictable behaviors that may happen. In principle, one would want to guarantee a certain level of robustness of the proposed technique against such uncertain phenomena. Stochastic optimization and randomized algorithms are possible ways of dealing with such uncertainty, both for designing new techniques, and for assessing the robustness of the available ones.

13.00 – 14.00 Keynote: Large-scale cluster management at Google with Borg
John Wilkes, Google, Mountain View

Google's Borg system is a cluster manager that runs hundreds of thousands of jobs, from many thousands of different applications, across a number of clusters each with up to tens of thousands of machines. It achieves high utilization by combining admission control, efficient task-packing, over-commitment, and machine sharing with process-level performance isolation. It supports high-availability applications with runtime features that minimize fault-recovery time, and scheduling policies that reduce the probability of correlated failures. Borg simplifies life for its users by offering a declarative job specification language, name service integration, real-time job monitoring, and tools to analyze and simulate system behavior.

14.00 – 15.30 Discussion 8 (Panel): The Future Challenges for Cloud Control

Thomas Bohnert, ZHAW, Zürich
Joe Butler, Intel, Leixlip
Anne Holler, VMware, Palo Alto
Azimeh Sefidcon, Ericsson Research, Kista
Bruno Sinopoli, Carnegie Mellon University, Pittsburgh
John Wilkes, Google, Mountain View

Thursday, June 11

8.30 – 8.50 Zero-touch cloud service assurance enabled by analytics
Andreas Johnsson, Ericsson Research, Kista, Sweden

Recent advances in computing, networking, analytics and control lead us to believe in and envision an artificial Telecom Cloud operator that can reduce operational expenditures by providing zero-touch management and service assurance. In this presentation we briefly discuss this vision as well as ongoing research initiatives that are aimed as milestones in a larger research effort. The presentation will be followed up in the discussion on "Analytics for Cloud Control and Management".

8.50 – 9.05 Bringing Probablistic Guarantees to Autoscaling Algorithms
Ahmed Ali-ElDin, Umeå University

Recently, numerous autoscaling algorithms have been proposed in the literature to leverage the elastic capabilities of Cloud platforms. Many of these autoscaling algorithms provision resources based on the current and the future predicted system load. Most of the algorithms suggested in the literature lack any guarantees on the Quality-of-Service (QoS) achieved when using the algorithm. In this work, we introduce a statistical framework that provides probabilistic guarantees based on well known statistical models on the behaviour of an autoscaler. The framework aims to calculate the amount of over-provisioning required in order to guarantee an error percentage in the autoscaler's predictions that is lower than a user-defined threshold. To achieve this, the proposed method requires a training phase where an error model for the autoscalers' predictions is constructed. This error model is then used to set the amount of overprovisioning required in order to achieve a certain level of QoS. We discuss the different possible QoS guarantees that such a framework might provide. We prove that the only possible guarantees that can be provided for an autoscaling algorithm is related to the number of times the algorithm mispredicts rather than the volume of the mispredictions. Using multiple state-of-the-art autoscaling algorithms and various workloads, we demonstrate the applicability and generality of our approach.

9.05 – 9.20 Performance Anomalies in Systems at Scale: Detection, Diagnosis, and Research Trends
Olumuyiwa Ibidunmoye, Umeå University

Performance anomalies are a norm rather than exceptions in large-scale systems. In fact, system bottlenecks are leading causes of costly SLA violations and unplanned down-times. The problem of uncovering and understanding such performance problems in different system and application domains is well studied. To assess progress, research trends and identify open challenges, we have reviewed and presented major research contributions in a survey accepted at the ACM Computing Surveys. In this talk, we present major findings from the survey and conclude with a brief of current work.

9.20 – 9.35 On Control-Theoretical Approaches to Cloud Spot Pricing
Zheng Li, Lund University

Spot pricing has been considered as a significant supplement for building a full-fledged market economy for the Cloud ecosystem. However, unlike the static and straightforward pricing schemes of on-demand and reserved Cloud services, the market-driven mechanism for pricing spot services would be complicated for both implementation and understanding. We propose to use control-theoretical approaches to deal with different problems related to Cloud spot pricing. For example, we have tried to use the Predator-Prey model to reveal the backend information behind spot price variations, and using a feedback control to help determine bids for a Cloud spot service. The long-term goal would be a control-based mechanism for generating spot prices.

9.35 – 9.50 Dynamic Application Placement in the Telco-cloud
Amardeep Mehta, Umeå University

To meet the challenges of consistent performance, low communication latency, and user mobility, we foresee a telco cloud that incorporates public cloud infrastructures with mobile access networks compute resource augmented telecom nodes. We identify key resource management challenges for this telco cloud and develop system models for data center, networks, and mobile application, including user mobility. Based on these models, we define an application placement optimisation problem that incorporates aspects of network link capacity, desired user latency and user mobility, as well as data center resource utilisation and server provisioning costs. We propose an application placement and migration algorithm based on local search. The proposed algorithm is evaluated in a simulated telco cloud environment for three scenarios, user mobility, diurnal usage patterns on a university campus, and rapid application popularity with a strong locality pattern during a sporting event. Our evaluation demonstrates that the placement and migration algorithm improves application latency, resource utilisation, and as such demonstrates the viability of dynamic resource management for the telco cloud.

10.20 – 10.40 MITOSIS: distributed autonoMIc managemenT Of ServIce compoSitions
Giovanni Toffetti, ZHAW, Zürich

The MITOSIS project aims at providing autonomous control for cloud-based services as resilient and scalable functionalities deployed within the service components in a distributed fashion. The goal is on one side to avoid vendor dependency/lock-in when dealing with service QoS management functionalities (e.g., automated health-management, auto-scaling, monitoring), and on the other to embed these functionalities within the service so as to make them resilient and scalable with the service as it grows.

10.40 – 12.00 Discussion 9: Analytics for Cloud Control and Management
Rolf Stadler, KTH & Andreas Johnsson, Ericsson Research

Recent advances in computing and memory technologies, as well as scientific progress in online algorithms and scalable computing make it feasible to analyse large amounts of operational data in real-time. In this discussion session, we will address the questions of how cloud management tasks can benefit from such analytics capabilities and what the associated research challenges are. While offline analysis and prediction, e.g., based on event logs, have been studied in the past, we will give priority to real-time analytics and its possible application to resource management, quality assurance, anomaly detection, etc.

10.40 – 12.00 Discussion 10: Self- applications versus cloud-vendor management services, are they worth the effort?*
Andy Edmonds & Giovanni Toffetti, ZHAW, Zürich

Running cloud applications entails much more than running components in VMs or containers. Cloud apps need to be monitored, scaled, updated and repaired around the clock. These management functionalities are fundamental for the correct functioning of cloud apps across their entire life cycle. They can be offered as services by IaaS / PaaS or third party providers, but are generally costly and might result in vendor lock-in. Using OSS PaaS solutions or components and adapting them to one's needs is an option, but development costs have to be considered and additional resource costs for running management functionalities need to be factored in.

10.40 – 12.00 Discussion 11: Self-adaptive network architecture for InfiniBand based HPC clouds
Feroz Zahid & Ernst Gunnar Gran, Simula Research, Fornebu

The research on network optimization in InfiniBand (IB) networks has been evolved in several directions, e.g. increasing network utilization, fault-tolerance, congestion control, and energy-aware systems. However, for efficient HPC clouds based on IB, the optimization problem becomes both complex and multi-dimensional, while individually proposed solutions often yield contradictory management decisions. We believe that a holistic closed-loop control system is required to effectively incorporate multidimensional objective function in future IB systems. Based on control theory, a self-adaptive model for the IB subnet system, may help achieving better network utilization while effectively keeping user level SLAs in HPC clouds.

13.00 – 13.15 The Wallenberg Autonomous Systems Program (WASP)
Karl-Erik Årzén, Lund University

13.15 – 14.00 Concluding discussions
Erik Elmroth, Umeå University

Feedback on the 7th edition of the workshop
Discussion on future events