

# Program

## 2<sup>nd</sup> CleanSky Conference

### August 04<sup>th</sup> - 05<sup>th</sup>, 2016



Hosted by NTNU-Norwegian  
University of Science and  
Technology, Trondheim

**Organizing Committee:**

Prof. Yuming Jiang, NTNU (Chair)  
Dr. Otto J. Wittner, UNINETT (Co-Chair)  
Dr. David Palma, NTNU  
Mr. Besmir Tola, NTNU (CleanSky Fellow)  
Ms. Yordanos T. Woldeyohannes NTNU (CleanSky Fellow)



START TIME	AUG 04 (Thursday) SESSIONS
08:45	Conference Opening
09:00	Keynote <i>"CleanG: NFV Derived Opportunities To improve Cellular Network Protocols"</i> Prof. K.K. Ramakrishnan, UC Riverside, USA
09:45	Coffee Break (15 mins)
10:00	Keynote <i>"A Hands-on Tutorial on the OpenNetVM NFV Platform"</i> Assist. Prof. Timothy Wood, The George Washington University, USA
12:00	Lunch Break
13:00	Guest Talk <ul style="list-style-type: none"> <li><i>"SDN Solutions for Scale-Out Data Centers"</i> Dr. Muzzamil Aziz, GWDG</li> </ul>
13:30	Fellow Talk (20+10 mins): Work by/Presenter <ul style="list-style-type: none"> <li><i>"Grouping users for better cache performance"</i> Pengyuan Zhou-ESR1/Song Yang-ER2</li> </ul>
14:00	Coffee Break (15 mins)
14:15	3x Fellow Talks (20+10 mins): Work by/Presenter <ul style="list-style-type: none"> <li><i>"Latency-sensitive data allocation for cloud storage"</i> Song Yang-ER2/Pengyuan Zhou-ESR1</li> <li><i>"MiddleBox Broker (MBB): Enabling In-network functions in secured traffic "</i> Alessio Silvestro-ESR2/Nitinder Mohan-ESR8</li> <li><i>"Cloud, Edge and Fog Computing: the art of the trade-off"</i> Nitinder Mohan-ESR8,/Alessio Silvestro-ESR2</li> </ul>
15:45	Coffee Break (15 mins)
16:00	2x Fellow Talks (20+10 mins): Work by/Presenter <ul style="list-style-type: none"> <li><i>"Joint VNF Placement and Flow Routing Optimization"</i> Yordanos T. Woldeyohannes-ESR3/Abhinandan S. Prasad-ESR7</li> <li><i>"Improving End to End Energy Efficiency in Network Function Virtualization"</i> Abhinandan S. Prasad-ESR7/Yordanos T. Woldeyohannes-ESR3</li> </ul>
17:00	Guest Talk <i>"SINet: Software-Defined Intermittent Networking"</i> Dr. David Palma, NTNU
19:00	Social Event (Dinner at Egon Tarnet)



START TIME	AUG 05 (Friday) SESSIONS
09:00	Keynote <b><i>"Network Neutrality: An Engineering Perspective"</i></b> Prof. Vishal Misra, Columbia University
10:15	Coffee Break (15 mins)
10:30	2x Fellow Talks (20+10 mins): Work by/Presenter <ul style="list-style-type: none"> <li>• <i>"Policy Engine as a Service (PEaaS): An Approach to a Reliable Policy Management Framework in Cloud Computing Environments"</i> Faraz F. Moghaddam-ESR5/Sameer Kulkarni-ESR4</li> <li>• <i>"DRENCH: Distributed Load Balancing for NFV based Service Function Chaining"</i> Sameer Kulkarni-ESR4/Faraz M. Moghaddam-ESR5</li> </ul>
11:30	Guest Talk <i>"Dependability requirements on the NFV Orchestrator and the SDN Controller"</i> Dr. Andres J. Gonzalez, Telenor
12:00	Lunch Break
13:00	4x Fellow Talks (20+10 mins): Work by/Presenter <ul style="list-style-type: none"> <li>• <i>"Cloud Service Reliability Assessment"</i> Besmir Tola-ESR9/Sachin Sharma-ER3</li> <li>• <i>"Enabling decomposition of a monolithic application into microservices and future steps"</i> Sachin Sharma-ER3/Besmir Tola-ESR9</li> <li>• <i>"Towards Virtual Private Cloud Embedding Through SDN in Multi-Tenants Cloud"</i> Konglin Zhu-ER1/Syed Shah-e-Mardan Ali Rizvi-ESR6</li> <li>• <i>"DVMP: Incremental Traffic-aware VM Placement on Heterogeneous Servers in Data Center"</i> Syed Shah-e-Mardan Ali Rizvi-ESR6/Konglin Zhu-ER1</li> </ul>
15:00	Coffee Break
16:00	Social Event Organized Visit at Nidaros Cathedral
19:00	Dinner (TBD)



## Keynote Biography

### ▀ Prof. K.K. Ramakrishnan, UC Riverside

K. K. Ramakrishnan is a Professor in the Dept. of Computer Science and Engineering at the University of California, Riverside. Until 2013, he was a Distinguished Member of Technical Staff at AT&T Labs Research, in Florham Park, New Jersey. Previously, K. K. was Technical Director and Consulting Engineer in the Networks Business Unit at Digital Equipment Corporation. K. K. has published over 200 papers and has been issued 151 patents in the United States. K. K. is an IEEE Fellow and an AT&T Fellow, recognized for contributions to the area of congestion control and virtual private networks. Between 2000 and 2002, he was at TeraOptic Networks, Inc., as Founder and Vice President. K. K.'s work on the "DECbit" congestion avoidance mechanism was recognized in 1995 retrospective issue of ACM Sigcomm Computer Communication Review as one of the 16 most important papers published over the previous 25 years in ACM Sigcomm publications. The work once again received the ACM Sigcomm "Test of Time" Paper Award in 2006. K. K. also received a Technology Medal from AT&T in 2013 for contributions to their Mobile Video Delivery Strategy and Optimization.

### ▀ Ass. Prof. Timothy Wood, The George Washington University

Timothy Wood is an assistant professor in the Department of Computer Science at George Washington University. Before joining GW, he received a doctoral degree in computer science from the University of Massachusetts Amherst and a bachelor's degree in electrical and computer engineering from Rutgers University. His research studies how cloud computing platforms can be built from massive data centers containing thousands of servers and storage devices. He seeks to improve the performance, reliability, and energy efficiency of these large distributed systems by adding automation and intelligence at the operating system and virtualization layers. His PhD thesis received the UMass CS Outstanding Dissertation Award, his students have voted him CS Professor of the Year, and he has won three best paper awards, a Google Faculty Research Award, and an NSF Career award.

### ▀ Prof. Vishal Misra, Columbia University

Vishal Misra is a Professor in the Department of Computer Science at Columbia University, with a joint appointment in the Electrical Engineering Department and an IEEE Fellow. His research emphasis is on mathematical modeling of networking systems, bridging the gap between practice and analysis. He has received the NSF and DoE Career awards as well as IBM and Google Faculty awards. He served as the Vice-Chair of the Computer Science Department at Columbia University from 2009 to 2011, and in 2011 he spun out Infinio, a company in the area of datacenter storage. He is also credited with inventing live-microblogging at Cricinfo, company he co-founded while a graduate student at UMass Amherst, predating Twitter by 10 years. Cricinfo was later acquired by ESPN and is still the world's most popular sports portal



## Detailed Program/Abstract

- ▀ Prof. K.K. Ramakrishnan, UC Riverside

### ***“CleanG: NFV Derived Opportunities To improve Cellular Network Protocols”***

Cellular Networks include complex protocols for mobility and mobile device management. Some of this can be attributed to the division of responsibility across multiple purpose-specific components in the packet core (aka EPC). NFV offers an opportunity to consolidate these components into a software environment, a direction that researchers and industry are taking. However, a significant opportunity to substantially simplify these protocols, rather than just re-implementing the functionality in software. This talk will discuss some of these opportunities.

- ▀ Ass. Prof. Timothy Wood, The George Washington University

### ***“A Hands-on Tutorial on the OpenNetVM NFV Platform”***

OpenNetVM is a high performance NFV platform based on Intel DPDK and Docker containers that supports packet processing at rates over 60 Gbps. This session will start with a brief presentation of the OpenNetVM architecture and a discussion of the resource management challenges found in NFV environments. It then will include a hands-on tutorial for using DPDK and OpenNetVM, including how to setup the platform, run network functions, manipulate service chains, and develop your own NFs. This will provide attendees a basic understanding of how to use our open source NFV platform for their own research on network middleboxes or NFV resource management.

- ▀ Prof. Vishal Misra, Columbia University

### ***“Network Neutrality: An Engineering Perspective”***

The issue of Network Neutrality has ignited considerable public debate recently. While the term and much of the discussion originated in the legal community, we started looking at it from an engineering and networking perspective a few years ago. We employed the lens of cooperative game theory and a careful modeling of the Internet including the topology, peering relationships and protocols used on the Internet. Our primary conclusion is that Network Neutrality should be expressed in terms of how you treat competition, not in how you treat packets and we proposed a definition of Network Neutrality that expresses that. We present some of our results including our prediction back in 2008 of a rise in paid peering (last year Netflix signed paid peering arrangements with all 4 of the top broadband providers in the US), the inadequacies of the Network Neutrality regulation in the US and the recent regulations in India, where the regulations are consistent with our definition of Network Neutrality.



- ▶ Dr. Muzzamil Aziz, GWDG

### ***“SDN Solutions for Scale-Out Data Centers”***

The current era of data center networks is facing a relentless growth of east-west traffic per year. Whereas, the traditional data center architectures are fundamentally designed for catering north-south traffic. To tackle such situations, NEPHELE EU project proposes an electronic-optical hybrid network architecture that scales linearly with the number of data center hosts and consolidates compute and storage networks over a single, Ethernet optical TDMA network. Moreover, a proposed SDN control plane framework of NEPHELE's data center architecture offers an easy monitoring and network management interface for future scale-out data center networks.

- ▶ Dr. David Palma, NTNU

### ***“SINet: Software-Defined Intermittent Networking”***

Operating in oceans and seas across the world are various types of vehicles, buoys and sensors. Moreover, their availability has increased in high-latitude regions, due to their socio-economic impact, creating several communication challenges. This presentation introduces the SINet project, which aims at developing an integrated communication system for intermittent links in the Arctic. Heterogeneity and networking robustness are two key aspects of this project, which proposes a Software-defined Intermittent Networking solution.

- ▶ Dr. Andres J. Gonzalez, Telenor

### ***“Dependability requirements on the NFV Orchestrator and the SDN Controller”***

This talk will focus on the elements that an NFV Orchestrator and an SDN controller need to consider in order to guarantee dependable operations. It describes the challenges posed from the standardization point of view, the internal dependability features of these two systems, and the technical requirements when interacting with the other different elements of the Network ecosystem.



- ESR 1: Pengyuan Zhou, Univ. of Helsinki

***“Grouping users for better cache performance”***

Cache strategy is an important factor of network performance because of its direct influence on helping content delivery, which improves the user experience. Researchers have proposed different cache strategies such as PropCache, Coordinating in-network caching, Probabilistic in-network caching etc. However, most of those strategies are based on content popularity. In this paper, we focus on identifying the benefit of grouping users according to content requests. More particularly, requests from users in a group will be routed to the same cache with an aim of improving cache hit ratio by exploiting the interest similarity within a group.

- ESR 2: Alessio Silvestro, NEC

***“MiddleBox Broker (MBB): Enabling In-network functions in secured traffic ”***

With increased adoption of HTTPS in the last several years, TLS is becoming the standard encryption protocol for internet connections and will become the default in upcoming years. TLS provides server authentication and connection encryption for end-to-end communications. However, TLS completely ignores application layer middleboxes which provided useful in-network services in HTTP thereby increasing the operation load and cost for network operators while gravely impacting the overall performance. Several solutions have been proposed for enabling in-network functions in TLS however, either such solutions do not provide control over these services, regarding how they use the data; or does not allow users to specify the dynamic security requirements. In this paper, the authors have proposed a platform named MiddleBox Broker (MBB) which provides an ad-hoc and user-controlled instantiation of middleboxes in network operator infrastructure. MBB provides several mechanisms to control and authenticate the in-network services offered by middleboxes. MBB logically separates an in-network service from an in-network instance such that the service providers can dynamically create instances based on user's workload. The authors also evaluate and discuss several aspects of MBB such as deployment methodologies, response delays, throughput and security.

- ESR 3: Yordanos T Woldeyohannes, NTNU

***“Joint VNF Placement and Flow Routing Optimization”***

Network Function Virtualization(NFV) envisions to bring flexibility in the management and deployment of network functions(NFs) such as firewalls, IDS, Proxy, etc also called middleboxes. NFV decouples the software of NFs from the physical hardware and run it on VM. This approaches will decrease the complexity in the daily network management and configuration operations observed now as well as decrease the OPEX and CAPEX. NFVs success strongly relies on the presence of VNF (Virtual Network Function) placement and routing algorithm that is able to decide how many instances to create in the network and where to place the instances. In this talk, a multi-stage and multi-objective VNF placement and flow routing model will be presented. The functionality of the model will also be demonstrated by using simulation results.



- ESR 4: Faraz Fatemi Moghaddam, GWGD

***“Policy Engine as a Service (PEaaS): An Approach to a Reliable Policy Management Framework in Cloud Computing Environments”***

Security Management in terms of validation of security requests, managing different protection levels for different customers with differing policies and capabilities of different service providers poses to be a major challenge and obstacle to the advancement of cloud computing and in particular for the IT-based on-demand services. This work PEaaS - Policy Engine as a Service, presents a reliable and standardized framework to address these security management challenges. PEaaS is a structural policy management engine that takes into account the capabilities of the cloud provider and the requirements of cloud customers to provide the dedicated security levels (rings). It enhances the process of policy management in cloud by performing syntactic and semantic analysis of the security requests based on the three-levels of protection ontology namely the protocol, mechanism and algorithm. PEaaS proposes an attribute based policy management model, and incorporates structural policy management engine, policy database, policy matchmaker, reasoning engine and an SLA engine, that enhance the management of security policies.

- ESR 5: Sameer Kulkarni, University of Gottingen

***“DRENCH: Distributed Load Balancing for NFV based Service Function Chaining”***

To enable the ability of dynamically instantiate network function instances (NFI), Drench has been introduced as a hybrid algorithmic framework for load balancing among virtual network function instances to combine distributed decision making with centralized control for orchestration and coordination while performing complex, dynamic service function chaining. There are several features such as ***shadow-price*** for aligning to the current resource availability, ***centralized SDN controller and OE (orchestration engine)*** for performing market orchestration and coordination activities, and ***novel distributed approach*** for dynamic service instantiation. To evaluate the performance of DRENCH, a prototype was developed based on large scale simulation and according to the results, DRENCH is able to quickly react to burst and elastic traffic by creating instances and redirecting while ensuring stability. Furthermore, the standard deviation of the loads across the NFIs is kept to a minimum by quickly instantiating services and redirecting flows and resources are efficiently utilized by timely termination of NFIs when they are not required at light loads.



- ESR 6: Syed Shah-e-Mardan Ali Rizvi, Tsinghua University

### ***“Towards Virtual Private Cloud Embedding Through SDN in Multi-Tenants Cloud”***

As the tremendous momentum cloud computing has grown, the modern data center networks are facing challenge to handle the increasing traffic demand among virtual machines (VMs). Simply adding more switches and links may increase network capacity but at the same time increase the complexity and infrastructure cost. Thus, intelligent VM placement has been proposed to reduce the intra-DC traffic. Prior solutions model the traffic-aware VM placement problem as a Balanced Minimum K-cut Problem (BMKP). However, the assumptions of “once-for all” VM placement on physical servers with equal VM slots are often not realistic in practical data centers, and thus the naive BMKP model may lead to suboptimal placement solutions. In this work, we revisit the problem by considering the server heterogeneity and propose an incremental traffic-aware VM placement algorithm. Given that the BMKP model cannot be directly applied, we make a number of transformations to re-establish the model. First, by introducing pseudo VM slots on physical servers with less VM slots, we allow the number of available VM slots of each server to be different. Second, pseudo edges with infinite costs are added between existing VMs, and thus previously deployed VMs on the same physical server will still be packed together. Third, a change on the number of pseudo VM slots is applied, so that existing VMs placed on different physical servers will still be separated. In this way, we reduce the problem to a new BMKP problem, which results in a much better solution. The evaluation results show that DVMP can reduce up to 28%, 39% and 55% traffic compared with naive BMKP model, greedy VM placement and random VM placement, respectively.

- ESR 7: Abhinandan S. Prasad, University of Gottingen

### ***“Improving End to End Energy Efficiency in Network Function Virtualization”***

Network Function Virtualization (NFV) is a new paradigm of virtualizing the services provided by the telecom and cloud operators. Given one or more service chains deployed on the virtual machines, the goal is to find a solution which minimize the power consumption and at the same time enforce service level agreement. We model the service chain as a robust queue and apply dynamic frequency and voltage scaling (DVFS). Our experiments prove the effectiveness of DVFS with respect to performance.



- ESR 8: Nitinder Mohan, University of Helsinki

***“Cloud, Edge and Fog Computing: the art of the trade-off”***

Fog computing is extending the Cloud Computing paradigm pushing the computation to the Edge of the network. Several application use-cases benefit from such paradigm such as Internet of Things (IoT), Smart Grid, Smart Cities and Wireless Sensors. However, different use-cases expose different requirements in terms of communication latency, number of nodes, nodes mobility, etc. This talk presents a novel algorithm for jobs placement on a Fog-Edge Network. It considers application-specific requirement (e.g., type of jobs, inter-job dependence, etc.) as well as Fog-Edge Network characteristics (e.g., number of nodes, link capacity, nodes' processing power, etc.). Moreover, this talk presents several evaluation tests performed on a Fog-Edge network simulator showing the impact of the architectures, the algorithms and the use-cases requirements on the network cost and processing time.

- ESR 9: Besmir Tola, NTNU

***“Cloud Service Reliability Assessment”***

Cloud-based applications and infrastructures enable us to conduct e-commerce transactions, interact with others via social applications, track our fitness, and even helps us to organize and plan our daily lives. Given the growth and our increasing reliance on cloud services, it is crucial that we have a good understanding of the reliability of these services. Moreover, today, a cloud computing service is composed of a complex array of building components and any of these can fail, which makes it a hard scenario to assess. Prior related studies on cloud reliability analysis have mostly focused on a limited set of cloud services and a few building blocks, and fails to capture a broader view of the reliability of cloud services. We leverage the fact that, today, cloud services are publicly exposing detailed information about failure incidents on the web. By crawling such publicly available incident reports, we conducted a large-scale measurement study of reliability of cloud services. Our analysis provides a better understanding of: the failure rate of a diverse set of cloud services and how the rate evolves over time, the different types of failures, recovery time for failure incidents, and availability of the services. We further investigate different factors potentially contributing to the disparity in reliability between services.



- ER 1: Konglin Zhu, Tsinghua University

***“The Embedding of Virtual Network to Physical Network in Data Centers”***

Network virtualization in the cloud allows multiple heterogeneous virtual cloud to coexist on a shared data center. Efficient embedding of the virtual cloud to the substrate data center is the first step toward the virtualization. The state-of-the-arts allow the network nodes and links mapping from virtual network to physical network separately. However, due to the limitation of the network technology, node mapping can only achieved in many-to-one manner. With the help of SDN, in this work, we propose a network embedding algorithm by considering the node splitting during the mapping process. The designed algorithm can support the dynamic change of network. We conduct the theoretical analysis to show the NP-hardness of the problem and compare the proposed algorithms with the states-of-the-arts to show the outstanding performance of the proposed algorithm.

- ER 2: Song Yang, GDWG

***“Latency-sensitive data allocation for cloud storage”***

Customers often suffer from the variability of data access time in cloud storage service, caused by network congestion, load dynamics, etc. One solution to guarantee a reliable latency-sensitive service is to issue requests with multiple download/upload sessions, accessing the required data (replicas) stored in one or more servers. In order to minimize storage costs, how to optimally allocate data in minimum number of servers without violating latency guarantees remains to be a crucial issue to tackle. In this talk, we study the latency-sensitive data allocation problem for cloud storage. We model the data access time as a given distribution whose Cumulative Density Function (CDF) is known, and prove that this problem is NP-hard. To solve it, we propose both exact Integer Nonlinear Program (INLP) and Tabu Search-based heuristic. The proposed algorithms are evaluated in terms of the number of used servers, storage utilization and throughput utilization.

- ER 3: Sachin Sharma, NEC

***“Enabling decomposition of a monolithic application into microservices and future steps”***

Microservice architecture is a method of deploying a large complex application as small, modular, and loosely coupled services in which each service runs independently and communicates through well-defined, lightweight mechanisms (such as REST API). We investigate microservice architectures and perform research on how these architectures can be adapted to be suited for future communications. In this talk, we research on when to start/stop decomposition of an application into microservices and discuss some of future research directions for complexity, configuration, communication, latency, and availability issues of microservices.

## List of Attendees

### CleanSky ITN Attendees

Name	Surname	Institution
Sameer	Kulkarni	University of Goettingen
Abhinandan	S Prasad	University of Goettingen
Fabian	Schneider	NEC Laboratories Europe
Alessio	Silvestro	NEC Laboratories Europe
Sachin	Sharma	NEC Laboratories Europe
Song	Yang	GWDG
Faraz	Fatemi Moghaddam	GWDG
Konglin	Zhu	Tsinghua University
Syed Shah-e-Mardan Ali	Rizvi	Tsinghua University
Yordanos Tibebe	Woldeyohannes	NTNU
Besmir	Tola	NTNU
Yuming	Jiang	NTNU
Otto J	Wittner	UNINETT
Pengyuan	Zhou	University of Helsinki
Nitinder	Mohan	University of Helsinki
K. K.	Ramakrishnan	University of California, Riverside

## List of Attendees

### External Attendees

Name	Surname	Institution
Steinar	Bjørnstad	TransPacket AS
Aryan	TaheriMonfared	University of Stavanger
Tesfaye Amare	Zerihun	NTNU
David	Palma	ITEM, NTNU
Poul	Heegaard	NTNU
Ameen	Chilwan	UNINETT
Min	Xie	Telenor
Zhifei	Mao	ITEM, NTNU
Gianfranco	Nencioni	NTNU
Peder J.	Emstad	NTNU
Håkon	Gunleifsen	NTNU Gjøvik
Ruxandra F.	Olimid	NTNU
Fengyou	Sun	ITEM, NTNU
Ilir	Bytyci	NTNU
Andres	Gonzalez	Telenor
Jing	Ma	NTNU
Timothy	Wood	The George Washington University
Vishal	Misra	Columbia University
Katina	Krlevska	ITEM, NTNU
Ramtin	Aryan	Høgskolen i Oslo og Akershus
Finn Arve	Agesen	NTNU

## Photos

- ▶ Prof. Yuming Jiang ( NTNU ) chairing the Conference



- ▶ Keynote of Prof. K. K. Ramakrishnan (Univ. California, USA)



## Photos

- ▶ Keynote of Ass. Professor Tim Wood (Univ. George Washington, USA)

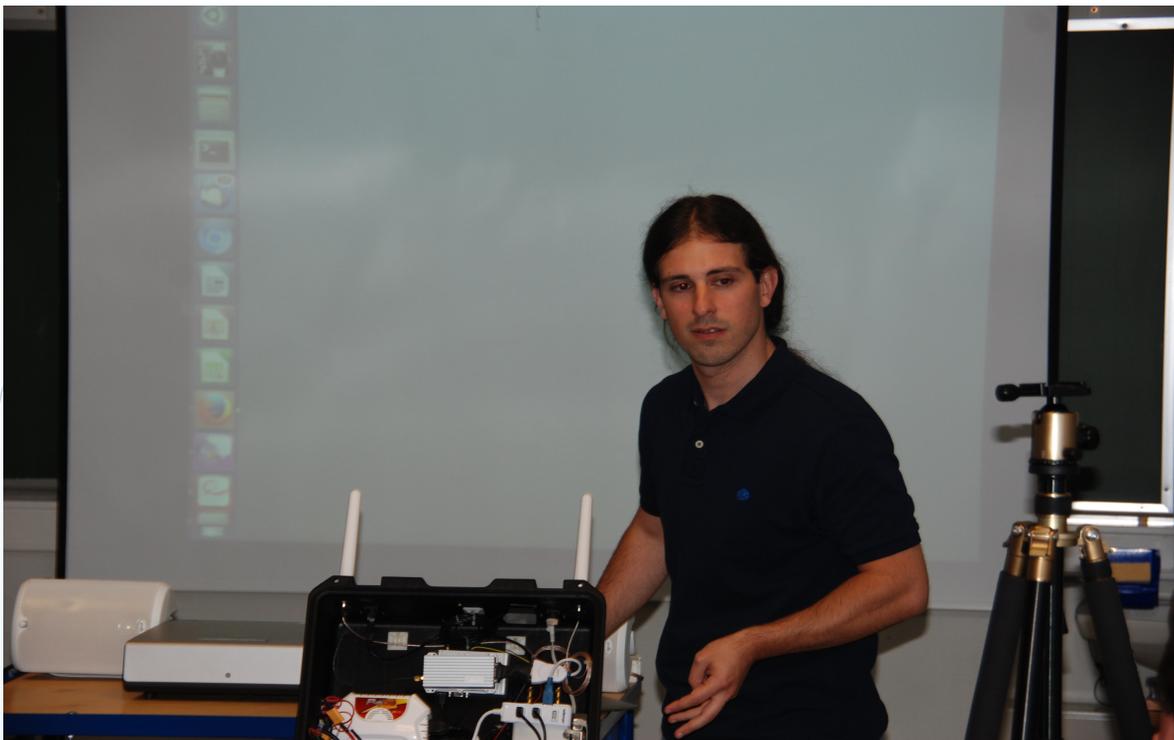


- ▶ Keynote of Prof. Vishal Misra (Columbia Univ., USA)



## Photos

- ▶ Guest talk and demo illustration of Dr. David Palma (NTNU, Norway)



- ▶ During the Lunch break

