\Orchestrating a brighter world NEC

Program
1st CleanSky Conference
February 29th & March 1st

Hosted by NEC Laboratories Europe in Heidelberg, Germany

Useful information & Directions

Social Event

- Monday 7pm @ Kulturbrauerei; Address: Leyergasse 6, 69117 Heidelberg http://www.heidelberger-kulturbrauerei.de/en/
- Dirk will walk from the Lab to the Brewery starting 18:30, feel free to join

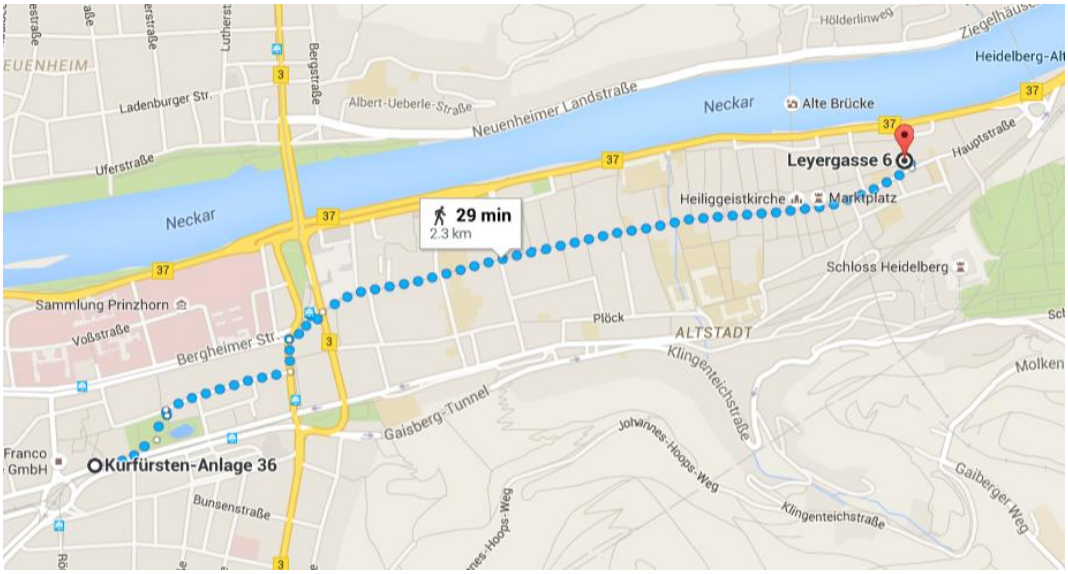

Agenda Day 1: Mon, Feb 29th

Time	Session
9:00	Keynote (50+10min) Distributed Data Structures over ICN Christian Tschudin, Uni Basel
10:00	3x Fellow Talks (15+5min) • Grouping users for better cache performance *Pengyuan Zhou, U Helsinki* • The Embedding of Virtual Network to Physical Network in DCs *Konglin Zhu, Tsinghua University* • An efficient service-chain forwarding mechanism in NFV *Syed Shah-e-Mardan Ali Rizvi, Tsinghua University*
11:00	Coffee Break (30min)
11:30	3x Fellow Talks (15+5min) • Secure and user-controlled in-network functions *Alessio Silvestro, NEC Laboratories Europe* • Distributed load balancer for NFV based SFCs *Sameer Kulkarni, Uni Göttingen* • Simplifying IoT computation in the cloud *Nitinder Mohan, Uni Helsinki*
12:30	Lunch (in-house, 90min)
14:00	3x Guest Talks (25+5min) • Middlebox Measurements and Cooperation *Mirja Kühlewind, ETH Zürich* • Software-Defined Multicast: Revisiting Multicast as an SDN-based Network Service for ISPs • *Jeremias Blendin & Julius Rückert, TU Darmstadt* • Impact of Heterogeneous Hardware Platforms on the Performance of Network Functions *Thomas Zinner, Uni Würzburg*
15:30	Coffee Break (30min)
16:00	3x Guest Talks (25+5min) • Kernel-Accelerated Packet Processing, *Oliver Hohlfeld, RWTH Aachen* • Enabling Communications in Disaster Scenarios with ICN *Jan Seedorf, FH Stuttgart* • Information Resilience through User-Assisted Caching in Disruptive Content-Centric Networks *Yiannis Psaras, UCL*
17:30	End of technical program
19:00	Social Event – Dinner @ Kulturbrauerei

CleanSky ITN

Agenda Day 2: Tue, Mar 1st

Time	Session
9:00	Keynote (50+10min) Back-Office traffic on the Internet Anja Feldmann, TU Berlin
10:00	2x Fellow Talks (15+5min) • Policy Management Engine (PME) to Improve the Reliability of Managing Resources in Clouds *Faraz Fatemi Moghaddam, GWDG* • Service composition effects on SLAs and nowadays Cloud reliability *Besmir Tola, NTNU*
10:40	Coffee Break (30min)
11:10	2x Fellow Talks (15+5min) • Resource Management in Datacenters and future challenges *Yordanos Tibebu, NTNU* • Reliable Virtual Machine Placement and Routing in Distributed Clouds *Song Yang, GWDG* Guest Talk (30+10min) • SONATA Johannes Lessmann, NLE
12:30	Lunch (in-house, 90min)
14:00	Panel (5x6min+60min) – Chair: Johannes Lessmann, NEC Labs Europe • Martin Stiemerling, FH Darmstadt & NLE • Philipp Wieder, GWDG • Anja Feldmann, TU Berlin • Jussi Kangasharju, Uni Helsinki • Ernö Kovacs, NLE
15:30	Coffee Break (30min)
16:00	2x Guest Talks (25+5min) • Using SDN-technologies to restore trust in network elements, *Martin Stiemerling, FH Darmstadt* • OpenState: Programming Platform-independent Stateful OpenFlow Applications Inside the Switch, *Marco Bonola, Uni Roma*
17:00	End of technical program

CleanSky ITN

Detailed Program

Abstracts

Orchestrating a brighter world NEC

Keynotes

Mon 9am:

“Distributed Data Structures over ICN” by Christian Tschudin, Uni Basel

While Information Centric Networking and its name-to-data binding is often discussed at a network and transport level, we posit that ultimately it will lead to a new way of handling distribute data structures. Starting from the primordial statement that ICN is about connecting data instead of hosts, we literally use names to link an ICN's data items. Using the core "lookup(name)" operation we will examine high-level data structures where names relate to pointers that permit to form complex objects e.g., index tables for file-like data collections. We also add some more speculative thoughts about the consequences for the patterns that names should expose, from the usefulness, or not, of version tags, to chunk numbers, catalogs and complete (functional) expressions.

Tue 9am

“Back-Office traffic on the Internet” by Anja Feldmann, TU Berlin

Most often when people think about the Web or the Internet they only consider the interactions between the browser and the front-end server. Indeed, although traffic between Web servers and Web browsers is readily apparent to many knowledgeable end users, fewer are aware of the extent of server-to-server Web traffic carried over the public Internet. We refer to the former class of traffic as front-office Internet Web traffic and the latter as back-office Internet Web traffic (or just front-office and back-office traffic, for short). Back-office traffic, which may or may not be triggered by end-user activity, is essential for today’s Web as it supports a number of popular but complex Web services including large-scale content delivery, social networking, indexing, searching, advertising, and proxy services. This talk takes a first look at back-office traffic, measuring it from various vantage points, including from within ISPs, IXPs, and CDNs. We then discuss implications and opportunities that the presence of back-office traffic presents for the Internet ecosystem.

CleanSky Fellow Session, Mon 10am

Grouping users for better cache performance

Pengyuan Zhou, U Helsinki

ICN has been seen as one of the most popular next generation network solutions. Cache deploy-ment is an importance topic of ICN research because of its direct influence on content delivery, which is a major metric of user experience. In this paper, we focus on the influence of varying user interest similarity on ICN cache system performance. Based on grouping users according to interest similarity, the requests of groups are directed to specified cache servers. We try to find the mapping relationship between cache system performance and interest similarity of inter and intra groups.

The Embedding of Virtual Network to Physical Network in Data Centers

Konglin Zhu, Tsinghua University

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 leave the theoretical analysis and evaluation as the future work for next stage.

An efficient service-chain forwarding mechanism in NFV

Syed Shah-e-Mardan Ali Rizvi, Tsinghua University

Future internet architecture has been advanced by such strong practices such as NFV (Network Function Virtualization) and SDN (Software Define Network). The research shows that if we combine their benefits, we can get much higher performance results in different cases such as cost effectiveness, network management and orchestration, effective cloud computations, data centeric developments and so much more related to the Future Internet Infrastructure. These advancements play a vital role in providing different kind of services using same physical plat-form which is very dominant application of these technologies. In such virtualized environment it becomes challenging to manage those services having so many virtualized forwarding nodes and to bind them into chains related to specific services. It would be the main interest to pursue future research that how those service chains should be handled by virtualized forwarding node.

Secure and user-controlled in-network functions

Alessio Silvestro, NLE

A significant part of the Internet Web traffic is currently encrypted using HTTPS. TLS, proposed by IETF in 1999 in the RFC2246, is the standard encryption protocol in HTTPS and will be most likely the *de facto* protocol used for all the traffic in HTTP/2. The main goal of TLS is to provide privacy and data integrity between two end hosts in the network. That is, it assumes that all the functionalities must reside at the endpoints of the communication, making it impossible to use in-network services wiping out the collaboration between Service Providers and Internet Ser-vice Providers. In-network functions, often called middleboxes, are used to optimize network resources, to improve user experience and to increase security. Introducing in-network func-tions into TLS sessions, is possible today only doing tricks that in the end will compromise the overall security. This talk will provide a brief introduction to the "Secure and user-controlled in-network functions" concept. This approach allows the deployment of in-network functions without com-promising important security aspects and enabling new scenario of cooperation between the various stakeholders (i.e., ISP, SP and users).

Distributed load balancer for NFV based Service Function Chains

Sameer Kulkarni, Uni Göttingen

In this work, I'll present our work 'DRENCH' – a semi-distributed SDN based framework that supports elastic scaling of NF instances and ensures load balancing across Network functions and present on the on-going and future research directions in terms of providing security and meeting the IoT use cases. The background and motivation of our work being: Middleboxes have become an integral part of Internet infrastructure, providing additional flow processing for policy control, security, and performance optimization. Network Function Virtualization (NFV) proposes the deployment of software based middleboxes on top of commercial off-the-shelf (COTS), enabling the dynamic adjustment of Virtual Network Functions (VNFs), both in terms of instance numbers and computational power. With logically centralized controller that has a global view of the network, SDN provides more flexibility to handle heterogeneous policies with multiple service chains and to steer traffic across network. However, the performance of Data center and Enterprise networks depend strongly on efficient scaling of VNFs and the traffic load balance across VNF instances, hence it is vital to optimally place the network functions, scale out and scale in the instances to meet the traffic demands and as well to adjust the load across different network functions.

Simplifying IoT computation in the cloud

Nitinder Mohan, Uni Helsinki

In a traditional cloud model, the data is pushed by the generators to a "central" cloud which processes and archives that data. The results are later retrieved by a querying client who often lies in the same layer as that of generators. The geographical location of the central cloud may be extremely far from both clients and generators, and this distance is bridged by underlying routers and gateways. In use-cases such as Internet-of-Things (IoT), the network fails to accommodate data generated by billions of sensors which ends up overwhelming the cloud. The networking issues directly affect the computation, mobility and security on IoT data. To refrain from sending bulk data to a central entity for computation, we propose a hybrid edge/fog cloud model that brings the computation closer to sensors and clients. The edge/fog layer resides at 1-2 hop distance from sensors and clients and provides computational capability through in-layer distribution. Processing data closer to IoT devices decouples the computation from network restrictions. The central cloud is responsible for providing data archival of the generated data. We decouple the architecture from single vendor deployments and provide security, mobility and processing priorities based on classification of the data.

Guest Talk Session, Mon 2pm

Middlebox Measurements and Cooperation

Mirja Kühlewind, ETH Zürich

Network operators increasingly rely on the use of in-network functionality provided by middleboxes to make their networks manageable and economically viable. These middleboxes need to make assumptions about the traffic passing through them which leads to ossification of the Internet protocol stack: new features or protocols are difficult to deploy because middleboxes don't understand them. The first step in fixing this situation is to gather data about the nature and distribution of middlebox impairments in the Internet. Further, based on this data we aim to develop a new middlebox cooperation protocol to explicitly communicate with these middleboxes to resolve their dependency on implicit assumption that might or might not be true. Both of these tasks, path transparency measurements and a new architecture for middlebox cooperation, are in focus of the H2020 MAMI project that just started this January.

Software-Defined Multicast: Revisiting Multicast as an SDNbased Network Service for ISPs

Jeremias Blendin & Julius Rückert, TU Darmstadt

Today's Internet is dominated by over-the-top (OTT) services with video accounting for the largest traffic share. Many of these services could greatly benefit from new network services. One such service is network-layer packet duplication, which IP multicast failed to provide due to a number of conceptual shortcomings. As a result, live video delivery today relies on huge CDN infrastructures to scale to Internet-wide audiences. However, they usually end at the edge of broadband access ISPs and use traffic-intensive IP unicast towards the end users. As a result ISPs are cut out as active stakeholder and the revenue chain, leaving no incentive to offer novel network services in an important part of the Internet. The advent of SDN fosters a fresh view on the natural domain-based structure of the Internet and enables new approaches for the introduction of new network services driven by the ISPs. In this talk, we present the concept of Software-Defined Multicast (SDM) as representative for a new kind of ISP-based network services. SDM and its extensions presented in this talk enable a seamless live media distribution with multicast efficiency using unicast packet semantics, while respecting the interests of all involved stakeholders in the OTT delivery process.

Impact of Heterogeneous Hardware Platforms on the Performance of Network Functions

Thomas Zinner, Uni Würzburg

The talk highlights two examples for the impact of heterogeneous hardware platforms on the performance of network functions. Firstly, the performance of a prototypical implementation of a mobile Serving Gateway (SGW) is investigated. The evaluations focus on the data plane performance of a pure software-based implementation and the corresponding DPDK-enabled implementation. Secondly, the efficient usage of heterogeneous switching hardware is presented. The software TableVisor enables the split of composite data path functions among several switches. Thus, advantages of specific switches can be used efficiently. TableVisor acts as a proxy and can thus be integrated easily in available SDN-enabled networks..

Guest Talk Session, Mon 3:30pm

Kernel-Accelerated Packet Processing

Oliver Hohlfeld, RWTH Aachen

As network speed increases, servers struggle to serve all requests directed at them. This challenge is rooted in processing overheads that mainly originate from the split between the kernel space network stack and user space applications. To address this challenge, we propose Santa, an architecture that allows user-space server applications to offload replies to common requests into the kernel. Our evaluation of HTTP and DNS as widely used protocols highlights that Santa substantially improves the packet processing performance of servers by up to 150% and up to 340%, respectively. Thus, Santa unlocks the speed of kernel-space networking for legacy server software without requiring extensive changes or specialized implementations.

Enabling Communications in Disaster Scenarios with ICN

Jan Seedorf, FH Stuttgart

Information Centric Networking (ICN) is a new paradigm where the network provides users with named content, instead of communication channels between hosts. The ICN approach can provider better security, data availability and network performance by leveraging data object security, in-network caching and per-hop forwarding mechanisms. The GreenICN project has developed a prototype for a future 5G network with a focus on decentralized operation and resilience against disruption. This talk will describe specifically how GreenICN leverages ICN features in a disaster scenario, e.g., at the aftermath of a hurricane or tsunami, when energy and communication resources are at a premium and it is critical to efficiently distribute disaster notification and critical rescue information. Key to this is the ability to exploit fragmented networks with only intermittent connectivity.

Information Resilience through User-Assisted Caching in Disruptive Content-Centric Networks

Ioannis Psaras, UCL

We investigate an information-resilience scheme in the context of Content-Centric Networks (CCN) for the retrieval of content in disruptive, fragmented networks cases. To resolve and fetch content when the origin is not available due to fragmentation, we exploit content cached both in in-network caches and in end-users' devices. Initially, we present the required modifications in the CCN architecture to support the proposed resilience scheme. We also present the family of policies that enable the retrieval of cached content and we derive an analytical expression/lower bound of the probability that an information item will disappear from the network (be absorbed) and the time to absorption when the origin of the item is not reachable. Extensive simulations indicate that the proposed resilience scheme is a valid tool for the retrieval of cached content in disruptive scenarios, since it allows the retrieval of content for a long period after the fragmentation of the network and the "disappearance" of the content origin.

Policy Management Engine (PME) to Improve the Reliability of Managing Resources in Clouds

Faraz Fatemi Moghaddam, GWDG

Despite the considerable benefits of cloud computing as an emerging technology, there are some resource management concerns such as managing access policies for allocated resources that have affected the reliability of cloud-based environments. The most challenging issue in this area is to provide an appropriate level of security for the stored data in cloud storages. In fact, each individual customer needs to be granted reliable security level(s) based on defined details in SLA. These security levels might be common for all customers or independent based on the data sensitivity. Applying a single security level for all stored data is not efficient and takes considerable processing power of scheduler to manipulate sensitive and also non-sensitive data. On the other hand, managing multiple security levels of resources is the most challenging concern in multi-level policy models and needs an appropriate and efficient algorithm. The most popular approach to express high-level security constraints is based on the usage of metadata and languages in scheduler for the specification of security policies. Accordingly, the main aim of this project is to propose an attribute based policy management model to enhance the reliability of managing resources in clouds.

Service composition effects on SLAs and nowadays Cloud reliability

Besmir Tola, NTNU

Service composition effects on SLAs:
Today cloud services have seen a tremendous growth as a consequence of the main advantages of Cloud Computing. This is accompanied with an increasing complexity in composing services to satisfy users requests for more and more complex services, as well as providing performance monitoring and service level comparisons. That is why an effective modeling of SLAs and contracts that consider the underpinning aggregation of different services composing the users requested service, is needed. In fact, nowadays SLAs lack in detailed expressivity in SLA specification for managing SLAs and contract composition makes it necessary that service providers should select appropriate and optimal services to be combined together in providing composite complex services. This SC problem, in cloud environments, such that both functional and non-functional (QoS) requirements are satisfied is a well known NP hard problem. We will take a look at the latest techniques used by so far to tackle this complexity.
Cloud reliability:
Many applications are moving into the cloud and as a consequence there is a lot of work in avoiding, detecting and repairing failures in order to build a more reliable cloud infrastructure and more reliable cloud applications. We are aware that many Cloud Providers call them selves as tipple/quadruple Providers, but to what extent this is true. Data is available today and many providers expose their real time status and we have the chance to collect data on cloud failures from different point of views, applications and Cloud Providers. We are considering collecting reliability data to achieve answers to different question like , how reliable are cloud services today, what are the different types of reliability issues, how long does it take for recovery etc. This way we can gain a valuable insight about the state of the cloud reliability today.

CleanSky Fellow Session, Tue 11:10am

Resource Management in Datacenters and future challenges

Yordanos Tibebu, NTNU

Efficient resource management is important for guaranteeing QoS and increasing profit. The heterogeneity of workloads and resource types in datacenters makes resource management a challenging task. In this talk, an overview of current approaches taken to address this issue will be presented. In addition, the role multi-objective optimization can play in resource allocation will be discussed. And challenges that need to be tackled, especially on the context of placement and routing of Virtual Network Function (VNFs) will be pointed out.

Reliable Virtual Machine Placement and Routing in Distributed Clouds

Song Yang, GWDG

In nowadays cloud computing systems, leveraging the virtualization technology, the customer's requested data computing or storing service is accommodated by a set of mutual communicated Virtual Machines (VM) in a scalable and elastic manner. These VMs are placed in one or more datacenter nodes according to nodes' capacities, failure probabilities, etc. The VM placement service reliability refers to the probability that at least one set of the whole customer's requested VMs operates during the entire requested lifetime. The placed VMs should obey the agreed-upon reliability, otherwise the service provider may face revenue loss.

In this talk, we first study the problem of placing at most H sets of requested k VMs in a network to minimize the total allocated capacity, such that the VM placement reliability requirement is satisfied and each VM pair has a communication delay no greater than the specified. We prove that this problem is NP-hard. We subsequently propose an exact Integer Nonlinear Program (INLP) as well as an efficient heuristic to solve this problem. We conduct simulations to compare the proposed algorithms with two existing heuristics in terms of acceptance ratio, average used storage and running time. Finally, we study the related reliable routing problem of establishing a connection over at most w link-disjoint paths from a source to a destination, such that the connection availability requirement is satisfied and each path delay is no greater than D. We show that this problem is NP-hard and devise an exact algorithm for it.

Guest Talks, Tuesday, 3:30pm

SONATA

Johannes Lessmann, NEC Laboratories Europe

The EU Sonata project is currently developing an open-source implementation supporting two basic areas: on the one hand a development environment for network service design and testing, on the other hand a service execution platform for automated service lifecycle management (incl. onboarding, deployment, orchestration and termination). This talk will give a basic introduction to service programming and execution in the NFV world, and also elaborate on Sonata's architecture design and several open issues that have to be addressed in this context.

Using SDN-technologies to restore trust in network elements

Martin Stiemerling, FH Darmstadt

Today's life depends on a number of infrastructure technologies, such as tap water supply, electricity, freight logistics to carry food and other goods, newspapers to inform the society, the Internet, etc. Even as all of them might look like independent resources, they are already sharing and will be even more sharing one common aspect: a data network to interconnect the computing systems in the back. Those computing systems are used, for instance, to order new food for a super market or to organize public gatherings. Typically, but not limited to, these systems are using the Internet to connect to each other. In turn, any network, from your home IoT network to the global Internet, needs networking components to make the service functional, e.g., these network elements are switches, routers, and home gateways. The basic assumption is that network elements perform their function as expected, i.e., free of software and hardware bugs, and also free of an unwanted functionality. This unwanted functionality can be, for instance, a tampered firmware by a intelligence apparatus, which allows to record all packets being handled by such a modified network element. This talk explores the possibilities to use techniques out of the space of Software Defined Networks (SDN) to reduce the primary attack surface for tampered network elements.

OpenState: Programming Platform-independent Stateful OpenFlow Applications Inside the Switch

Marco Bonola, Uni Roma 2 & CNIT

Software Defined Networking envisions smart centralized controllers governing the forwarding behavior of dumb low-cost switches. But are "dumb" switches an actual strategic choice, or (at least to some extent) are they a consequence of the lack of viable alternatives to OpenFlow as programmatic data plane forwarding interface? Indeed, some level of (programmable) control logic in the switches might be beneficial to offload logically centralized controllers (de facto complex distributed systems) from decisions just based on local states (versus network-wide knowledge), which could be handled at wire speed inside the device itself. Also, it would reduce the amount of flow processing tasks currently delegated to specialized middleboxes. The underlying challenge is: can we devise a stateful data plane programming abstraction (versus the stateless OpenFlow match/action table) which still entails high performance and remains consistent with the vendors' preference for closed platforms? We posit that a promising answer revolves around the usage of extended finite state machines, as an extension (super-set) of the OpenFlow match/action abstraction. We concretely turn our proposed abstraction into an actual table-based API, and, perhaps surprisingly, we show how it can be supported by (mostly) reusing core primitives already implemented in OpenFlow devices.

List of Attendees

CleanSky ITN Attendees

Faraz Fatemi Moghaddam	GWDG
Jussi Kangasharju	University of Helsinki
David Koll	University of Göttingen
Sameer Kulkarni	University of Göttingen
Dirk Kutscher	NEC Laboratories Europe
Nitinder Mohan	University of Helsinki
Fabian Schneider	NEC Laboratories Europe
Jan Seedorf	NEC Laboratories Europe
Syed Shah-e- Mardan Ali Rizvi	Tsinghua University
Alessio Silvestro	NEC Laboratories Europe
Yordanos Tibebu Woldeyohannes	NTNU Trondheim
Besmir Tola	NTNU Trondheim
Philipp Wieder	GWDG
Song Yang	GWDG
Pengyuan Zhou	University of Helsinki
Konglin Zhu	Tsinghua University

External Attendees

Jeremias Blendin	TU Darmstadt
Marco Bonola	Uni Roma 2
Anja Feldmann	TU Berlin
Oliver Hohlfeld	RWTH Aachen
Ernö Kovacs	NLE
Mirja Kühlewind	ETH Zürich
Johannes Lessmann	NLE
Ioannis Psaras	University College London
Julius Rückert	TU Darmstadt
Martin Stiemerling	FH Darmstadt / NLE
Christian Tschudi	Universität Basel
Thomas Zinner	Universität Würzburg

Fellow Presentation Slides

Presentation slides available for download

Please feel free to download the slides presented by CleanSky ITN fellows. Simply follow this link.

Session notes available for download

During each session, the CleanSky ITN fellows took notes of the material presented and the questions asked in the Q&A sections. These notes are also available for download. Simply follow this link.

Photos

Keynote of Prof. Christian Tschudin, Universität Basel

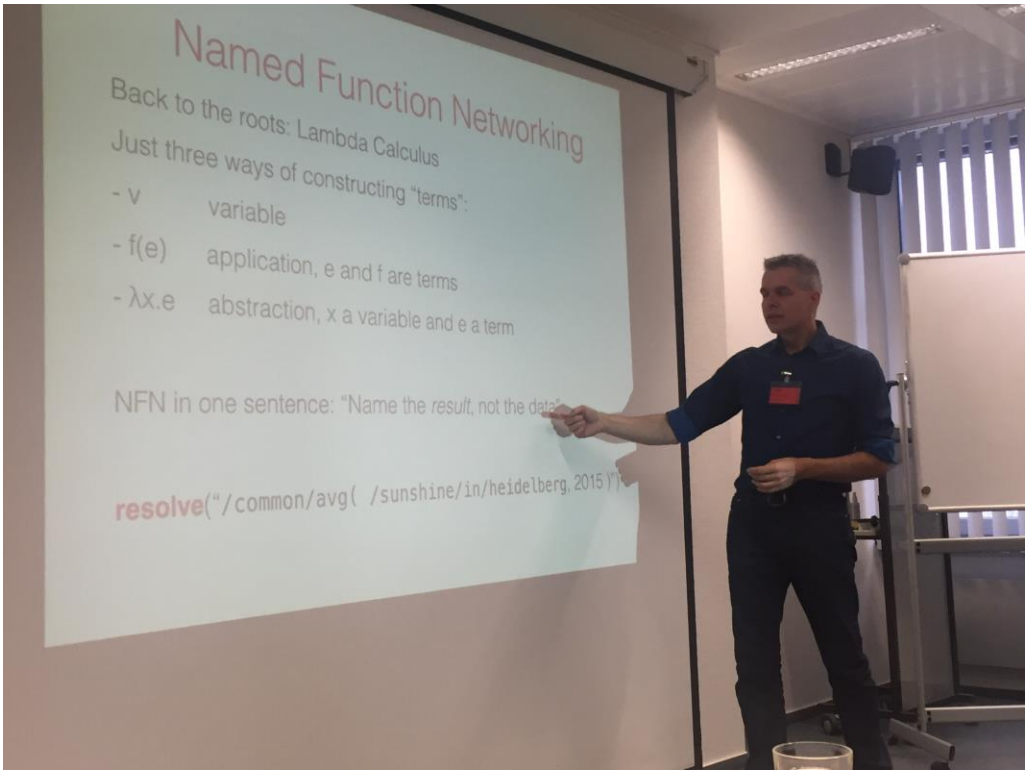

Guest Talk by Mirja Kühlewind, ETH Zürich

Photos

Social Event: Dinner

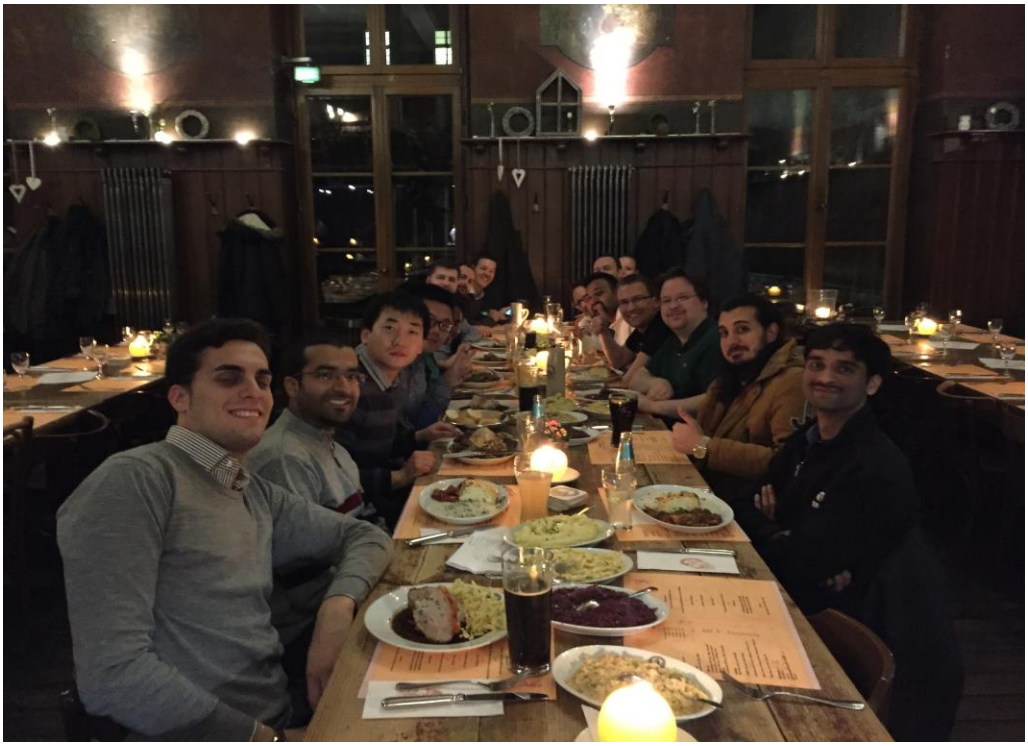

\Orchestrating a brighter world

NEC

NEC brings together and integrates technology and expertise to create the ICT-enabled society of tomorrow.
We collaborate closely with partners and customers around the world, orchestrating each project to ensure all its parts are fine-tuned to local needs.

Every day, our innovative solutions for society contribute to greater safety, security, efficiency and equality, and enable people to live brighter lives.