A Report on the Munich Internet Research Retreat 2017

Vaibhav Bajpai TU Munich bajpaiv@in.tum.de

ABSTRACT

This article summarizes the two-day Munich Internet Research Retreat (MIR) held in May 2017. The goal of the retreat was to provide a forum for both academic and industrial researchers to exchange ideas and get feedback on their current work. It was organized in a spirit that is similar to highly interactive "Dagstuhl" seminars, with a very limited number of full-length talks, while dedicating most of the time to poster sessions and group discussions. Presentations delivered during the seminar are made publicly available [20].

1. INTRODUCTION

The MIR originated from informal discussions of different research groups at TUM and a team at NetApp on diverse topics related to networking. The discussions brought together PhD students and post-docs to present their respective research (including both work in progress as well as polished results) and provided an informal setting for intense and rich exchange among participants involved. We realised that there was notable potential in reaching out further, which eventually led to the instantiation of the MIR.

The main mission of the MIR is to ensure mutual awareness of different teams working on current (complementary) topics in networking. We want to lay the foundations for establishing, broadening, and deepening cooperation among a variety of groups doing networking research. In order to foster easily sustainable relationships, our initial scope has been deliberately limited to the area around Munich (which may reach as far as 400 km in some cases). As a common denominator, we target like-minded teams within the region, where the common mindset stems from practical research in networked systems, paired with interest and efforts in the Internet Engineering Task Force (IETF), the Internet Research Task Force (IRTF) and the ACM SIGCOMM and SIGMOBILE communities.

The purpose of the MIR is threefold: 1) We seek to provide recurring opportunities for companies to get in touch with research groups that have expertise in fields relevant to the former. 2) We aim to support researchers in understanding current and emerging research and engineering problems from the commercial development and deployment perspectives. 3) We like to offer reality feedback to academic researchers and out-of-the-box ideas to those from industry. Overall, we hope to foster future bi- or multi-lateral collaboration between academics and industry.

The retreat is organized in a highly interactive fashion, combining posters (for providing variety) and group discussions intertwined with plenary talks that stimulate discussions. Organization directions are shaped by the feedback of the participants, keeping the format constantly evolving. We borrow some elements from the renowned Dagstuhl seminars: We limit the number of participants to ~ 40 to maintain interactivity and allow all participants to meet one another. We hold the retreat in Raitenhaslach away from the daily activities to ensure focus and include an overnight stay and a social dinner to foster continued interaction and allow for digesting ideas. The seminar is by invitation only, and we put an emphasis on the industry, picking PhD students with matching topics, which helps with obtaining a compatible and energetic mix of people. Because we know that everybody's time is scarce, we organize each retreat in a way that it occupies just two days including arrival and departure. With a target of two workshops per year, presently scheduled for May and November, we shall be able to continuously engage with a growing regional community even if individuals cannot participate on every occasion.

Towards this mission, the 1st MIR retreat was organized on November 24–25, 2016 at the TU Munich (TUM) Science and Study Center inbb Raitenhaslach, Germany. A 2nd iteration of the MIR was organized at the same location and held on May 23–24, 2017. Presentations on topics such as: ... were solicited. The retreat consisted of six invited presentations and several posters presenting early and upcoming research, with several breakout sessions to discuss topics of interests in an informal setting. Synopses of these sessions are described in this report in more detail.

2. INVITED PRESENTATIONS

The invited presentations were intended as a basis for triggering discussions and identifying areas for group work.

2.1 Opportunistic Content Dissemination

Many of the existing opportunistic networking systems have been designed assuming a small number links per node and have trouble scaling to large numbers of potential concurrent communication partners. In the real world we often find wireless local area networks with large numbers of connected users – in particular in open Wi-Fi networks provided by cities, airports, conferences and other venues. In this talk, Teemu Kärkkäinen (TU Munich) presented a 50 client opportunistic network in a single Wi-Fi access point and use it to uncover scaling problems and to suggest mechanisms to improve the performance of single segment dissemination. Further, we present an algorithm for breaking down a single dense segment dissemination problem into multiple smaller but identical problems by exploiting resource (e.g., Wi-Fi channel) diversity, and validate our approach via simulations and testbed experiments. The ability to scale to high density network segments creates new, realistic use cases for opportunistic networking applications.

2.2 User Tracking Based on TLS CCA

The design and implementation of cryptographic systems offer many subtle pitfalls. One such pitfall is that cryptography may create unique identifiers potentially usable to repeatedly and precisely re-identify and hence track users. Quirin Scheitle (Technical University of Munich) presented his investigation of TLS Client Certificate Authentication (CCA), which currently transmits certificates in plain text. He demonstrated [29] CCA's impact on client traceability using Apple's Apple Push Notification service (APNs) as an example. APNs is used by all Apple products, employs plain-text CCA, and aims to be constantly connected to its backend. Its novel combination of large device count, constant connections, device proximity to users and unique client certificates provides for precise client traceability. He shows that passive eavesdropping allows to precisely re-identify and track users and that only ten interception points are required to track more than 80 percent of APNs users due to global routing characteristics. The work was conducted under strong ethical guidelines, with responsibly disclosing the findings, and a working patch by Apple for the highlighted issue was confirmed. The aim for this work is to provide the necessary factual and quantified evidence about negative implications of plain-text CCA to boost deployment of encrypted CCA as in TLS 1.3.

2.3 Dynamic MultiPath Routing Protocol

Safety critical communication platforms often deploy multiple heterogeneous wireless link layer access technologies like ETSI TETRA, IEEE 802.11, IEEE 802.15.4, 3GPP LTE, satellite terminals or proprietary waveforms. Depending on the interface characteristics, vendor specific decisions and other aspects, these often come shipped with suitable mobile ad-hoc routing protocols to form networks in an autonomous manner independently for each link type. 802.11 links typically deploy OLSR, BATMAN or 802.11s, whereas proprietary waveforms often come with proprietary MANET protocols. To bridge these access networks at a logically higher level and provide an opaque operational network view, a network routing protocol at the top is required. Exterior routing protocols like BGP are limited in their use and features like automatic neighbor detection or reduced message overhead are required. Dynamic Multipath Routing (DMPR) tries to address these shortcomings and provides exterior routing protocol features for heterogeneous link layer environments even in low bandwidth environments. Furthermore, DMPR features policy based routing to route traffic through different paths if required or advantageous.

2.4 IoT Security: TrustZone for v8-M

Internet of Things (IoT) devices today use microcontrollers that are limited in CPU performance, as well as in RAM and flash size. Many of these devices use the ARM A-class or M-class processors. A-class CPUs are able to run popular operating systems (OS), such as embedded Linux, while M-class CPUs use Real-Time Operating Systems. So far, the security functionality of M- and A-class CPUs has been very different. A-class CPUs can use sophisticated hardware features, such as TrustZone offering physical separation between the normal and the secure world operating systems. M-class CPUs offer basic security protection using the memory protection unit (MPU), which offers memory isolation. Many IoT products, however, use very few operating system security techniques, if they run an OS at all, and often do not make use of hardware security support. While the exact reasons for these design decisions remain unclear, the growing list of IoT security failures calls for improved protection capabilities. With the introduction of the TrustZone for ARMv8-M architecture, security features known from the mobile world are now available in the IoT environment. Additionally, ARM v8-A processors enhance security with the Pointer Authentication Extension, which prevents return-oriented programming-based attacks. In this talk Hannes will explain the hardware security features offered by upcoming ARM processors and microcontrollers.

2.5 Redesigning Stack, API and Networks

Emerging Non-Volatile Main Memories (NVMMs), also known as storage-class memory and persistent memory push the majority of end-to-end latency that includes durable I/O to network stacks and their APIs. This not only impairs inherent performance of NVMMs that is one to two orders of magnitude faster than traditional persistent medias like Solid-state Drives (SSDs), but prevents systems from adopting them to be reliable with relative ease. Michio Honda (NEC) investigates towards solving this problem, designing an efficient network stack and its APIs, and exploring new opportunities in networking such as software switches and middleboxes in addition to improving networked storage systems.

2.6 State of Linux Network Development

Florian Westphal (Redhat) gave a brief summary of recent and ongoing development work in the linux kernel network stack, such as NIC hardware offloads (switching, routing, IPsec, TLS) and eXpress Data Path (XDP). After an overview of enhancements of linux TCP stack, such as BBR congestion control [7] and RACK loss detection [8] current development efforts were presented. This includes an extension of the BSD socket API to provide a full zero-copy interface, making qdiscs lockless and a planned removal of the IPv6 Forwarding Information Base (FIB) reader-writer lock.

3. PARALLEL GROUP WORK

The afternoon sessions were used to discuss selected topics in more depth in smaller groups. This section summarizes the discussions of each group.

3.1 Use Cases of and Research on Blockchains and other Distributed Ledgers

Blockchains gained a lot of attention since the raise of the Bitcoin [21] payment system. In essence, Blockchains can be understood as distributed database systems that offer their users unmodifiable and immutable storage for various types of application data. Today, Blockchains are typically used in the context of payment systems.

In this breakout session, the participants discussed how Blockchain technology can be used to increase accountability in various scenarios: 1) The first scenario is network administration. Instead of pushing new configurations to devices, administrators might deposit configurations for devices in a Blockchain-like system. The configurations are then polled and applied automatically by devices. This approach would ensure that configuration changes by, for instance, rogue administrators trying to harm their employer could be detected. The outlined system could also be expanded with a feature that allows for multi-party approval of configurations to prevent hazardous (or erroneous) configurations. 2) Another scenario would be accountability for log data from autonomous systems. Examples include "black box"-data of autonomous cars or drones. A trusted element might send a stream of status infromation to some kind of data center which uses Blockchain technology to persist the received information.

Besides use cases for Blockchain-like systems, the participants discussed available technical implementations. One difficulty is that "mining"-based (Proof-of-Work) [2] implementations cannot be used in (private) networks with modestly small numbers (\sim 50–100) of Blockchain nodes [13]. Hence, different approaches to reach consensus in the network need to be employed. The participants exchanged their knowledge about existing implementations that are suitable in private networks. In the opinion of the participants, the Hyperledger system [17] is one of the most promising contenders.

3.2 Measurements and Reproducibility

Computer science and engineering community has traditionally been less oriented towards providing reproducible research [4, 26] and more about quickly publishing novel and promising results. In this group, we discussed the current state and possible directions for improving the state of reproducibility.

The discussion began with questions regarding user privacy and sharing sensitive (network traffic, user identities, location, et al.) data since both legal and ethical concerns often inhibit data sharing. The first step that must be taken is data anonymization, which is difficult, sometimes computationally demanding and furthermore, it may only work for a specific research question. Once the data is sufficiently cleared of sensitive information and prepared for publication, long-term planning for storing large amounts of data is necessary. Additionally, even when there is willingness to take the necessary steps and enrich the research community by providing collected data, the risk of a certain party's intentional de-anonymization always remains. Actions that can be taken to enable more reproducible research were also discussed. For instance, Burkhart [6] explores cryptographic alternatives to data anonymization. Specifically, he revisits applying secure multiparty computation (MPC) to the problem of aggregating network data from multiple domains and develops a new framework with MPC operations for processing high volume data in near real-time.

Without the detailed methodology, much of a reproduction efforts remain guesswork. A detailed description is essential since it serves to validate others' work; not providing enough information about the methods used, for example regarding data processing, can cast a serious doubt on the methodology and any of the results. Provision of all the steps of the applied methodology also assists other practitioners and enables faster progress of research.

From the author's perspective, there is a disbalance between the efforts to prepare reproducible results and the potential reward. Since publications with novel measurement data and/or methodologies are likely to get more citations, for many authors the citation count is seen as a strong enough incentive. Reproducers find it even less rewarding since replicating others' results is time-consuming and receives little recognition. Although many conference call for papers express the need by explicitly asking for papers that reproduce previous work, there is currently little acceptance of reproduction papers: most of the studies – if ever materialize into a submission – get filtered out in the peer review round (due to the lack of novelty). Furthermore, rebuttal papers on previously published results are often subjected to rejection.

Instead of using very specific datasets for examining algorithm and protocol performance, having benchmark datasets could be one part of the solution. However, in this case the question that emerged was how we actually build suitable datasets for different applications and use-cases. The other action item would be to promote awareness that the current practices do not completely follow the scientific approach. This may be achieved by organizing more workshops with the focus primarily on reproducible research and validation of previous results, or by tweaking the reviewing process to encourage more openness.

We also discussed what can we learn from other fields where validation of results through replication has been an essential component. One of the examples with very high requirements for validation is in physics, where numerous teams work simultaneously on the same dataset and evaluate core models from many angles. However, this field deals with zero privacy and human subjects issues [3].

3.3 P4 and SDN

P4 [5] is a language that allows to program the structural layout of protocol headers, as well as, processing operations performed on those. As use cases discussed in this breakout session, P4 as extension for OpenFlow, which is limited by the fields that can be matched on, was discussed first. The second sample discussed was P4 when used for packet inspection on a switch. Complex protocol stacks (e.g. VXLAN) can be parsed on ingress, de-capsulating the then to be inspected payload, before the packet is de-parsed on the egress path again. Thereby, headers can be dropped or new ones added. The second topic discussed was possible extensions to P4, which are required to perform common data-center operations. The prime example is scheduling, required for time sensitive protocols (e.g. various IEEE 802.1 standards), as well as offering Quality of Service guarantees. Lastly, a brief discussion of recent activity [1, 10, 27] in the field of P4 was held, resulting in the exchange of new interesting papers.

3.4 IoT and security

Distributed Denial-of-Service Attack (DDoS) attacks involving IP-based cameras have led to discussions about the impact of large numbers of unpatched IoT devices. While the owners of the those devices may not see a need to update them, even if they are aware of their vulnerability, there is the risk of collateral damage for the Internet infrastructure as a whole. In the breakout session we collected concerns and listed challenges with IoT security in light of the possibility of such DDoS attacks.

The regulatory outreach possibility and whether firmware update functionality needs to be mandatory-to-support in an IoT product were disscussed. Since these firmware/software updates may, however, require user consent (or user action) the user needs to be contacted first. Definitely, user consent is required if the update does not only fix a security bug but is otherwise bundled with new features [22, 16] that may impact the privacy settings or the features of the product. How users can be contacted is an area for further investigation, considering that IoT devices may expose no user interface. Furthermore, the role network operators could play in detecting and stopping attacks was explored. Relationships with work on botnet mitigation [18, 9] by Comcast, a US cable operator, was drawn. In general, the idea of detecting attacks and mitigating them via network segmentation was attractive to various participants and the potential role of intermediaries, in the style of edge / fog computing, was discussed. This lead to further questions concerning the potential increase in attack surface posed by these entities and the changes in the security architecture due to their introduction into the communication path.

The group concluded the discussion with additional questions: What happens with those IoT devices that cannot be patched or where the support has expired? Should unpatchable IoT products be recalled? The car served as an example where a regulatory framework sets a requirement for recalls. As such, should there by an IoT TUeV (MoT in UK), a term borrowed from the car industry where vehicles are tested for roadworthiness in regular intervals? Who should pay for these services? What is the scope of the guarantee for consumers? Should devices have a kill switch to disable them from being connected to the Internet? Should there be a regulatory requirement for IoT device manufacturers, OEMs and service providers to inform users already at the time of purchase about the expiry date after which no further service will be available?

Overall, the a large part of the discussion reflected feelings of the participants that a regulatory framework will be necessary to provide incentive to avoid collateral damage. Technical measures would then follow.

3.5 Networking APIs

The first phase in the discussion of the group was to define the assumptions under which the flaws and possible improvements of the Networking APIs, in particular the standard Socket API, should take place. These were the following two basic assumptions:

- The currently deployed Internet architecture is assumed. New networking paradigms that would require completely new functionalities of the Networking-API are out-of-scope, e.g.: information centric / content based networking, delay-tolerant networking, vehicular networking.
- Specific requirements of **IoT are not considered** in the discussion since the IoT devices have very specific requirements, i.e. regarding energy-efficiency, so that the Socket API is often not applicable/is not being used. If Socket API is used in an IoT device, energy-consumption is not a problem since sending consumes much more energy than an inefficient implementation of the Socket API (e.g. copying data from userland to kernelspace).

Afterwards, the discussion focussed on problems observed with the current APIs. For high performance networking applications with, e.g., thousands of TCP connections, the API does not scale well. Several work-arounds have been developed in order to cope with this inefficiency, e.g. *sendfile* (directly copying from file descriptor to a socket without copying data to user-space) or *sendmessage* with zero-copying / zero-copy sockets (page re-mapping between kernel/user-space). The importance of hardware-offloading has increased. However, several open questions remain: What functions should be performed in hardware? Should the interface to the hardware be standardized?

The current API is fine for most standard cases but it becomes a bottleneck for high performance applications: This is mainly caused by the effort of copying data, often caused by a not packet-oriented processing of data in the application (i.e. application sends stream of data, transport layer builds segments). Again, workarounds have been developed, e.g. fast packet processing in userland or StackMap [31] + netmap framework [23] (dedicated NIC for one application, etc.).

This lead to a discussion of desirable properties of a networking API. Mentioned were the isolation of networking-stack and application, energy efficiency (mobile applications) and high performance and scalability (data center applications).

Possible solution ideas which were identified during the breakout were the application of dedicated I/O CPUs, integration of GPGPU processing and networking (offloading on

GPU), packetized processing of data in the application, and several techniques that could reduce the processing-overhead in kernel, e.g. avoid queuing of TCP ACK packets or reducing the overhead of a system call. Due to the limited time, a detailed discussion of the suitability of theses approaches would need to be done in a follow-up discussion. Furthermore, two talks on the second day of the retreat presented additional information on related subjects: Michio Honda (see Section 2.5) presented information on persistence in networking (redesigning stack, API and networks) and Florian Westphal (see Section 2.6) talked about current developments regarding the Linux networking stack.

4. POSTERS

Participants were encouraged to bring posters to present their recent research work.

4.1 Dynamic MultiPath Routing

Safety critical communication platforms often deploy multiple heterogeneous wireless link layer access technologies like ETSI TETRA, IEEE 802.11, IEEE 802.15.4, 3GPP LTE, satellite terminals or proprietary waveforms. Depending on the interface characteristics, vendor specific decisions and other aspects, these often come shipped with suitable mobile ad-hoc routing protocols to form networks in an autonomous manner independently for each link type. 802.11 links typically deploy OLSR, BATMAN or 802.11s, whereas proprietary waveforms often come with proprietary MANET protocols. To bridge these access networks at a logically higher level and provide an opaque operational network view, a network routing protocol at the top is required. Exterior routing protocols like BGP are limited in their use and features like automatic neighbor detection or reduced message overhead are required. DMPR tries to address these shortcomings and provides exterior routing protocol features for heterogeneous link layer environments even in low bandwidth environments. Furthermore, DMPR features policy based routing to route traffic through different paths if required or advantageous.

4.2 PASTE: A Networking Interface for NVMMs

Emerging Non-Volatile Main Memories (NVMMs), also known as storage-class memory and persistent memory push the majority of end-to-end latency that includes durable I/O to network stacks and their APIs. This not only impairs inherent performance of NVMMs that is one to two orders of magnitude faster than traditional persistent medias like SSDs, but prevents systems from adopting them to be reliable with relative ease. Michio Honda (NEC) presented an investigation of this problem, designing an efficient network stack [15] and its APIs, and exploring new opportunities in networking such as software switches and middleboxes in addition to improving networked storage systems.

4.3 Measuring the Performance of Mobile Users

In a mobile network, the mobile terminal (MT) continuously exchanges link related metrics and signals to the nearby base station to measure the strength and quality of the received signal. Quality of Service (QoS) metrics are used for handover decisions and cell reselection. A handover can occur if there is a strong radio signal in the neighboring cell while the serving cell's radio signal is getting diminished. However, previous studies [28] show that it is not always the value of signal strength that matters to have a good throughput performance. Therefore, knowing the possible achievable throughput value before making a handover is equally important along with link-related QoS metrics. Ermias Walelgne (Aalto University) proposes a solution to estimate the throughput value of posthandover using the metrics collected from the current serving base station. The result of this throughput prediction can be combined with other link QoS metrics such as RSSI and RSPQ values for better handover decision.

4.4 Lightweight Virtualization for Smart Cars

Modern vehicles are equipped with several interconnected sensors on board for monitoring and diagnosis purposes; their availability is a main driver for the development of novel applications in the smart vehicle domain. Roberto Morabito [19] presented a Docker container-based platform as solution for implementing customized smart car applications. Through a proof-of-concept prototype—developed on a Raspberry Pi3 board—we show that a container-based virtualization approach is not only viable but also effective and flexible in the management of several parallel processes running on On-Board Unit. More specifically, the platform can take priority-based decisions by handling multiple inputs, e.g., data from the CANbus based on the OBD II codes, video from the on-board webcam, and so on. Results are promising for the development of future in-vehicle virtualized platforms.

4.5 Data-driven Mobility Modeling

Ljubica Pajević Kärkkäinen (TU Munich) presents an analysis of a large trace of user associations in a university wireless network, which includes around one thousand access points on five campus sites. The trace was obtained from authentication logs of the RADIUS server collected over 16 months. User access patterns, specifically the arrival processes of users to wireless access points, visiting time duration, as well as the user arrival patterns to the buildings in the campus are studied. She observed that that a large fraction of the network - around half of all access points - exhibit Poisson arrival process, which is advantageous for modeling and prediction of network occupancy. By analyzing duration of associations with access points, she shows that the visiting time distributions can be modeled by two-stage hyper-exponential distributions. While network associations in campus wireless networks have been extensively studied in the literature, this study reveals changing access patterns, which seem to be characteristic for networks of predominantly mobile users.

4.6 iConfig - What I See is What I Configure

Michael Haus (TU Munich) presented iConfig to manage IoT devices in smart cities. The management of IoT devices in urban areas is becoming important due to that the majority of the people living in cities and the number of deployed IoT devices are increasing. Therefore, iConfig addresses three major issues in current IoT management: registration, configuration, and device maintenance. To achieve the goals of iConfig, the presented system relies on programmable edge modules, which can run on smartphones, wearables, and smart boards to configure physically proximate IoT devices.

4.7 Opportunistic Content Dissemination

Many of the existing opportunistic networking systems have been designed assuming a small number links per node and have trouble scaling to large numbers of potential concurrent communication partners. In the real world we often find wireless local area networks with large numbers of connected users - in particular in open Wi-Fi networks provided by cities, airports, conferences and other venues. In this talk, Teemu Kärkkäinen (TU Munich) presented a 50 client opportunistic network in a single Wi-Fi access point and use it to uncover scaling problems and to suggest mechanisms to improve the performance of single segment dissemination. Further, we present an algorithm for breaking down a single dense segment dissemination problem into multiple smaller but identical problems by exploiting resource (e.g., Wi-Fi channel) diversity, and validate our approach via simulations and testbed experiments. The ability to scale to high density network segments creates new, realistic use cases for opportunistic networking applications.

4.8 A Group Recommender System for Trips

Recommender systems (RSs) in tourism often recommend single Points of Interests (POIs) such as restaurants or museums. However, tourists visiting a destination are usually looking for a tourist trip composed of multiple POIs along a practical route. Daniel Herzog (TU Munich) presented a Recommender system (RS) [30] recommending tourist trips to a group of users. This is a particularly complex problem as the RS has to aggregate the travel preferences of all group members before generating recommendations. Furthermore, he wants to research on how different devices and user interfaces can support groups in providing feedback on recommendations and finding a consensus.

4.9 Data Dissemination in Vehicular Networks

Lars Wischhof (Hochschule Müncher) presented an architecture and preliminary results of an on-going research project at the research group where communication schemes combining cellular communication with direct-communication (such as Device-to-device (D2D) modes of the latest LTE-A releases or LTE-V) are combined for applications in intelligent mobility. The basic assumption is that future vehicles will most-likely have multiple communication technologies and modes available. Therefore, a context-aware selection of the communication mode is advocated. A suitable architecture is outlined. First simulation results for the example of a DENMbased application indicate that a context-aware selection can outperform a static assignment.

4.10 Accountability for Cyber-Physical Systems

Severin Kacianka (TU Munich) seeks to capture the essential features of an accountable (computer-)system. Logs are, for example, a common way to create evidence and establish "truth" in computer systems. Another facet are mechanisms to process those logs and techniques to formulate the questions of compliance with laws as queries against those logs. However, there are currently no "blue prints" on how to make a system "accountable". We wish to develop a comprehensive framework that makes it possible to explicate the accountability features of a system, reason about their effectiveness, compare it to other solutions and offer options to exchange one specific component for another.

4.11 Real-time TE in the Internet

Edwin Cordeiro (TU Munich) aims to create a network solution capable of detecting and avoiding congestions in the Internet. The avoidance is done automatically by a central controller adapting running network protocols and network routes to avoid congestion. Such objectives are divided in two parts. The first one is the creation of a method capable of detecting congestion in the Internet in real-time without probes at destinations. The second is the implementation of Software Defined Networks (SDN) ideas in the routing system using the Interface to the Routing System (I2RS) protocol, that is being specified by IETF.

4.12 Fine-Grained Edge Offloading for IoT

Vittorio Cozzolino (TU Munich) makes the case for IoT edge offloading, which strives to exploit the resources on edge computing devices by offloading fine-grained computation tasks from the cloud closer to the users and data generators (i.e., IoT devices). The key motive is to enhance performance, security and privacy for IoT services. The proposal bridges the gap between cloud computing and IoT by applying a divide and conquer approach over the multi-level (cloud, edge and IoT) information pipeline. To validate the design of IoT edge offloading, a unikernel-based prototype is developed and evaluated the system under various hardware and network conditions.

4.13 CARISSMA

Center of Automotive Research on Integrated Safety Systems and Measurement Area (CARISSMA) is a new center for vehicle safety. Its focus is on passive as well as active vehicle safety research. The main goal is to develop a global safety system to support 'Vision Zero', achieving the ultimate goal of zero traffic deaths. Therefore, all relevant disciplines are combined in one building. Nine professors alongside 47 staff members pursue are working in, e. g. an indoor driving area and full-vehicle crash test facility, a drop tower, a HiL- lab, a full-vehicle test bench, a lab for safe energy storage, a car2x-laboratory, a simulation lab and an open-air ground for performing full-vehicle tests.

4.14 Car2X Lab

Research on Car2X communication, the wireless communication between vehicles and other road participants, is gaining a major part of CARISSMA. The Car2X research facilities feature a powerful simulation computer, which helps to leverage our Open Source in-house Car2X simulation tools "Artery" and "Vanetza". These tools will also be integrated in our Car2X experimental vehicle, blending virtual simulation and real test drives for semi-virtualized testing. In addition to this, the research is focusing on HIL-testing of Car2X equipment. Further research topics are Teleoperated Driving and the use of Mobile Edge Computing for vehicular applications.

4.15 SENDATE

The goal of the SENDATE sub-project PLANETS is to provide cutting edge technical and scientific solutions for ProgrammabLe Architecture for distributed NETwork functions and Security. As TUM's Lehrstuhl für Kommunikationsnetze (LKN) is one of the main contributors of the SENDATE PLAN-ETS, this poster aims to present an overview of only some of the research directions of LKN. The most related work could be classified as i) Network virtualization evaluation, hypervisor control plane migration protocol (i.e. DITRA) and HyperFlex virtualization platform, ii) Resilient network design and planning & techno economic analysis and iii) Hypervisor placement, performance evaluation and mapping and synthetic IP graph generation.

4.16 Quantifying Flexibility in Networks

Communication networks need to face fast and frequent changes, since the request of resources from users is increasingly dynamic. Several new technologies have arisen to deal with this requirement, such as Software Defined Networking (SDN), Network Virtualization (NV) and Network Function Virtualization (NFV). These technologies claim to enhance the flexibility of the network, i.e., the ability of the network to adapt to changes. However, a formal definition of flexibility has not been settled, thus it is difficult to measure and compare the performance of these technologies. Alberto Martinez (TU Munich) proposed an initial definition of flexibility and we illustrate it with an use case: the dynamic placement of the controller in a SDN-enabled network.

4.17 Internet Architecture and Security

Quirin Scheitle (TU Munich) presented research directions in both Internet Architecture and Security questions. Examples for this are Geolocation of Routers using Ripe Atlas [24], mapping the communication flows and backend infrastructures of Mobiles Messaging Services such as WhatsApp or WeChat [25], and about creation and operation of an IPv6 hitlist [12].

4.18 Towards an Information Model for Decentralized Anomaly Detection

The DecADe project (Decentralized Anomaly Detection) [11] researches how anomaly detection can be implemented in the highly segmented and hierarchical networks of (autonomous) cars and aircraft cabins. Based upon the information model created in their previous work [14], Holger Kinkelin (TUM) outlined how data collected by software sensors deployed in the respective network segments can be collected and joined in socalled Forensics Centers. The unified view on the entire system provided by the Forensics Center allows for more comprehensive analyses compared to the restrained view of anomaly detection components that reside in a network segment.

5. CONCLUSIONS AND NEXT STEPS

The 2nd Munich Internet Research retreat concluded successfully on May 23–24, 2017. All the presentation material and contact information of presenters are available online [20]. The readers are encouraged to contact the organizers to learn more about the the next retreat.

We also collected some feedback from the participants.

Acknowledgements

This seminar was located at the TUM Science and Study Center in Raitenhaslach, Germany, supported by NetApp, Huawei, and TUM. The organizers would like to thank the participants (alphabetically ordered by first name) for their contributions-Alberto Martínez Alba (TUM) Alf Zugenmaier (Hochschule München) Arend Martin (BMW AG) Christian Facchi (TH Ingolstadt) Christoph Nufer (Rohde & Schwarz) Daniel Herzog (TUM) Dirk Kutscher (Huawei German Research Center) Dominik Scholz (TUM) Edwin Cordeiro (TUM) Ermias Walelgne (Aalto University) Florian Westphal (Red Hat) Georg Carle (TUM) Hagen Paul Pfeifer (Rohde & Schwarz) Hannes Tschofenig (ARM) Holger Kinkelin (TUM) Johannes Naab (TUM) Jörg Ott (TUM) Lars Eggert (NetApp) Lars Wischhof (Hochschule München) Ljubica Pajević Kärkkäinen (TUM) Marius Strobl (NetApp) Michael Haus (TUM) Michio Honda (NEC) Nemanja Deric (TUM) Quirin Scheitle (TUM) Roberto Morabito (Ericsson Research Finland) Severin Kacianka (TUM) Simon Leinen (SWITCH) Stefan Neumeier (TH Ingolstadt) Teemu Kärkkäinen (TUM) Vaibhav Bajpai (TUM) Vittorio Cozzolino (TUM)

6. **REFERENCES**

- A. Abhashkumar, J. Lee, J. Tourrilhes, S. Banerjee, W. Wu, J. Kang, and A. Akella. P5: policy-driven optimization of P4 pipeline. *Proceedings of the Symposium on SDN Research, SOSR 2017, Santa Clara, CA, USA, April 3-4, 2017*, pp. 136–142, 2017.
- [2] A. Back. Hashcash A Denial of Service Counter-Measure. Webpage, last visited 2017-05-31. URL: http: //www.hashcash.org/papers/hashcash.pdf.



- [3] F. Bai, N. Sadagopan, and A. Helmy. IMPORTANT: A framework to systematically analyze the impact of mobility on performance of routing protocols for adhoc networks. *Proc. IEEE Infocom*, pp. 825–835, 2003.
- [4] V. Bajpai, M. Kühlewind, J. Ott, J. Schönwälder, A. Sperotto, and B. Trammell. Challenges with reproducibility. *Proc. ACM SIGCOMM*, 2017.
- [5] P. Bosshart, D. Daly, G. Gibb, M. Izzard, N. McKeown, J. Rexford, C. Schlesinger, D. Talayco, A. Vahdat, G. Varghese, and D. Walker. P4: programming protocolindependent packet processors. *Computer Communication Review*:87–95, 2014.
- [6] M. Burkhart. *Enabling Collaborative Network Security with Privacy-Preserving Data Aggregation*. PhD thesis, 2011.
- [7] N. Cardwell, Y. Cheng, C. S. Gunn, S. H. Yeganeh, and V. Jacobson. BBR: congestion-based congestion control. *Commun. ACM*:58–66, 2017.
- [8] Y. Cheng, N. Cardwell, and N. Dukkipati. RACK: a timebased fast loss detection algorithm for TCP. Internet-Draft draft-ietf-tcpm-rack-01, Oct. 2016.
- [9] C. Chung, A. Kasyanov, J. Livingood, N. Mody, and B. V. Lieu. Comcast's Web Notification System Design. RFC 6108 (Informational), Internet Engineering Task Force, Feb. 2011. URL: http://www.ietf.org/rfc/rfc6108.txt.
- [10] H. T. Dang, H. Wang, T. Jepsen, G. J. Brebner, C. Kim, J. Rexford, R. Soulé, and H. Weatherspoon. Whippersnapper: A P4 language benchmark suite. *Proceedings* of the Symposium on SDN Research, SOSR 2017, Santa Clara, CA, USA, April 3-4, 2017, pp. 95–101, 2017.
- [11] DecADe BMBF Project. Webpage, last visited 2017-05-31, 2016. URL: https://www.net.in.tum.de/sites/decade/.
- [12] O. Gasser, Q. Scheitle, S. Gebhard, and G. Carle. Scanning the ipv6 internet: towards a comprehensive hitlist. *Proc. 8th Int. Workshop on Traffic Monitoring and Analysis*, 2016.

- [13] V. Gramoli. On the Danger of Private Blockchains (When PoW can be Harmful to Applications with Termination Requirements). Webpage, last visited 2017-05-31, 2016. URL: https://www.zurich.ibm.com/dccl/papers/ gramoli_dccl.pdf.
- [14] N. Herold, H. Kinkelin, and G. Carle. Collaborative Incident Handling Based on the Blackboard-Pattern. ACM Workshop on Information Sharing and Collaborative Security, pp. 25–34, 2016.
- [15] M. Honda, L. Eggert, and D. Santry. PASTE: Network Stacks Must Integrate with NVMM Abstractions. *Proc.* ACM HotNets, pp. 183–189, 2016.
- [16] HP detonates its timebomb: printers stop accepting third party ink en masse. http://boingboing.net/2016/09/19/ hp-detonates-its-timebomb-pri.html.
- [17] Hyperledger. Webpage, last visited 2017-05-31, 2017. URL: https://www.hyperledger.org.
- [18] J. Livingood, N. Mody, and M. O'Reirdan. Recommendations for the Remediation of Bots in ISP Networks. RFC 6561 (Informational), Internet Engineering Task Force, Mar. 2012. URL: http://www.ietf.org/rfc/rfc6561. txt.
- [19] R. Morabito, R. Petrolo, V. Loscri, N. Mitton, G. Ruggeri, and A. Molinaro. Lightweight Virtualization as Enabling Technology for Future Smart Cars. *International Symposium on Integrated Network Management* (*IM*), 2017.
- [20] Munich Internet Research Retreat 2017. Presentation Materials: https://www.cm.in.tum.de/en/mir/2017.
- [21] S. Nakamoto. Bitcoin: A Peer-to-Peer Electronic Cash System. Webpage, last visited 2017-05-31, 2008. URL: https://bitcoin.org/bitcoin.pdf.
- [22] Philips pushes lightbulb firmware update that locks out third-party bulbs. http://boingboing.net/2015/12/14/ philips-pushes-lightbulb-firmw.html.

- [23] L. Rizzo. Netmap: A novel framework for fast packet I/O.
 2012 USENIX Annual Technical Conference, Boston, MA, USA, June 13-15, 2012, pp. 101–112, 2012.
- [24] Q. Scheitle, O. Gasser, P. Sattler, and G. Carle. Hloc: hints-based geolocation leveraging multiple measurement frameworks. *Network Traffic Measurement and Analysis Conference (TMA)*, 2017.
- [25] Q. Scheitle, M. Wachs, J. Zirngibl, and G. Carle. Analyzing locality of mobile messaging traffic using the matador framework. *Passive and Active Measurements Conference (PAM) 2016*, Mar. 2016.
- [26] Q. Scheitle, M. Wählisch, O. Gasser, T. Schmidt, and G. Carle. Towards an ecosystem for reproducible research in computer networking. *Proc. ACM SIGCOMM*, 2017.
- [27] A. Sivaraman, S. Subramanian, M. Alizadeh, S. Chole, S. Chuang, A. Agrawal, H. Balakrishnan, T. Edsall, S. Katti, and N. McKeown. Programmable packet scheduling at line rate. *Proc. ACM SIGCOMM*, pp. 44–57, 2016.
- [28] S. Sonntag, L. Schulte, and J. Manner. Mobile network measurements - it's not all about signal strength. 2013 IEEE Wireless Communications and Networking Conference (WCNC), Shanghai, Shanghai, China, April 7-10, 2013, pp. 4624–4629, 2013.
- [29] M. Wachs, Q. Scheitle, and G. Carle. Push Away Your Privacy: Precise User Tracking Based on TLS Client Certificate Authentication. *Network Traffic Measurement* and Analysis Conference (TMA), Jun. 2017.
- [30] W. Wörndl, A. Hefele, and D. Herzog. Recommending a sequence of interesting places for tourist trips. *J. of IT* & *Tourism*:31–54, 2017.
- [31] K. Yasukata, M. Honda, D. Santry, and L. Eggert. Stackmap: low-latency networking with the OS stack and dedicated nics. 2016 USENIX Annual Technical Conference, USENIX ATC 2016, Denver, CO, USA, June 22-24, 2016. Pp. 43–56, 2016.