GI-Dagstuhl Seminar 16353 on Aware Machine to Machine Communication

Mayutan Arumaithurai, Stephan Sigg, Xiaoyan Wang

version: April 26, 2017

Abstract

This article summarizes a five day GI-Dagstuhl Seminar on 'Aware machine-to-machine communication' held from August 28th to September 2nd 2016. The seminar was a follow up of a DFG-JSPS seminar held in October 2015, in Karuizawa, Japan. The Karuizawa meeting focused on information-sharing in IoT, Crowd sensing and crowd steering, in-network data reduction, and Self-organizing data-collection networks. This second seminar focused on the aspects 'Security', 'Services', and 'Context' in Machine-to-Machine communication. It brought together people who are actively involved in the ICN community and also researchers with background on (usable) security, smart environments. The entire set of presentations delivered during the seminar is made publicly available at http://materials.dagstuhl.de/index.php?semnr=16353.

1 Executive Summary

Current trends show that machine-to-machine (M2M) interactions such as Internet of Things (IoT), wearables, vehicular networks and smart homes will play a major role in the Internet. In fact, it is expected that M2M interactions will constitute more than a third of the total connections[1]. These networks are rapidly growing in complexity and continuing to extend into the personal and private domain. Fuelled by the numerous sensors interconnected, massive amounts of data need to be managed and routed efficiently.

At the same time, networking technology is shifting towards virtualization, with Software Defined Networking (SDN) and Network Function Virtualization (NFV) likely to change the infrastructure landscape. The cloud concept is transforming the Internet to a network of data centers, with a communication model consisting of computer-to-cloud-to-computer interactions. Big-Data/Analytics based decision making is also expected to play a major role. Networking paradigms are witnessing a shift from the traditional end-to-end connections and location oriented networking to content/information oriented

[1] http://www.cisco.com/c/en/us/solutions/collateral/service-provider/ip-ngn-ip-next-generation-network/white_paper_c11-481360.html

architectures. Information Centric Networking (ICN), a popular future internet architecture that provides features such as the ability to handle content by its name, to secure individual pieces of data and support ubiquitous caching that allows data to be obtained from the closest source.

The potential benefits of combining this massive environmental perception based on M2M with the control power available in upcoming network paradigms is huge, as is the number of research issues opened. ICN and SDN have been primarily designed for fixed networks. While these technologies have the potential to cater to the needs to M2M based applications, there remains a lot of unresolved issues. This concerns, for instance, the largely unsolved question of scalability of ICN routing schemes, orchestration of NFV based services, as well as the location and actual implementation of SDN controllers. Some of the interesting research issues are listed below.

Context-based support for M2M The consideration of environmental stimuli for a flexible adaptation of networking and routing strategy can further advance current protocols. For instance, (1) Emergency situations might demand other routing schemes and priorities than normal operation; (2) DTN obtains a new dimension when the dissemination strategy can be changed according to, for instance, flow behaviours, movement speed or transportation mode; (3) Local networks could automatically be formed and secured against intruders in meeting situations or conditioned on social relation or friendship.

In M2M, where content covers also environmental situation and personal behaviour patterns, context-based services are capable of guiding towards not only content, but situation or groups of people with equal mindset, behaviour patterns or sentiment. Likewise, the location-independency of CCN-based routing holds significant potential of simplification for situation-dependent services: Instead of hard-coding individual sensor sources for input, always the nearest sensor reading of one particular type of source could be adaptively chosen.

Security for M2M Security and access control are key concerns for M2M and traditional end-to-end security approaches might not be sufficient to handle the plethora of use-cases envisioned for M2M. ICN, with its focus on securing individual pieces of content complimented with group encryption approaches such as ABE could be a potential solution for M2M security. SDN holds the promise of increased security when, for instance, security presets are conditioned on shared situation or also on friendship relationships. Friendship and social contact can then control security settings of each individual connection when all communication partners and their situation/context can be clearly identified or could associate themselves to ABE and use it for group encryption.

Services for M2M The massive amount of data available in such networks also demands for novel, efficient routing, storage and data reduction schemes. Moreover, Big-Data/Analysis based mapping and other decision making services might facilitate efficient M2M interactions. Publish/Subscribe services might

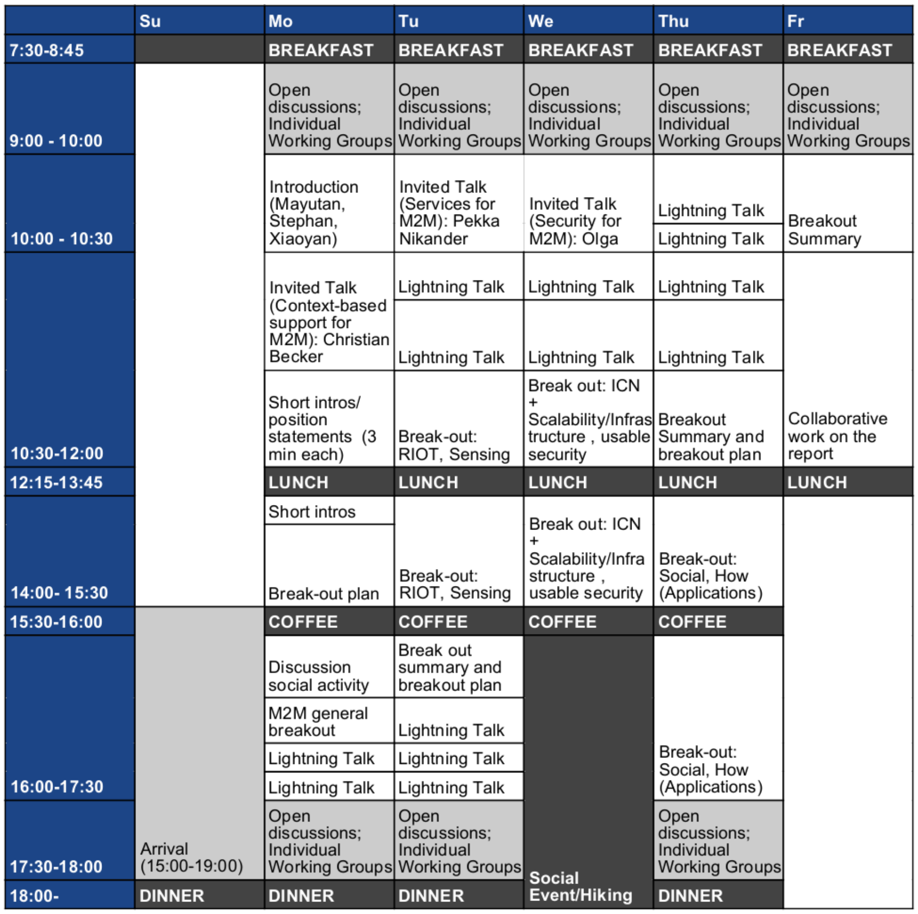

	Su	Mo	Tu	We	Thu	Fr
7:30-8:45		BREAKFAST	BREAKFAST	BREAKFAST	BREAKFAST	BREAKFAST
9:00 - 10:00		Open discussions; Individual Working Groups	Open discussions; Individual Working Groups	Open discussions; Individual Working Groups	Open discussions; Individual Working Groups	Open discussions; Individual Working Groups
10:00 - 10:30		Introduction (Mayutan, Stephan, Xiaoyan)	Invited Talk (Services for M2M): Pekka Nikander	Invited Talk (Security for M2M): Olga	Lightning Talk Lightning Talk	Breakout Summary
10:30-12:00		Invited Talk (Context-based support for M2M): Christian Becker Short intros/ position statements (3 min each)	Lightning Talk Lightning Talk Break-out: RIOT, Sensing	Lightning Talk Lightning Talk Break out: ICN + Scalability/Infrastructure , usable security	Lightning Talk Lightning Talk Breakout Summary and breakout plan	Collaborative work on the report
12:15-13:45		LUNCH	LUNCH	LUNCH	LUNCH	LUNCH
14:00- 15:30		Short intros Break-out plan	Break-out: RIOT, Sensing	Break out: ICN + Scalability/Infrastructure , usable security	Break-out: Social, How (Applications)	
15:30-16:00	Arrival (15:00-19:00)	COFFEE	COFFEE	COFFEE	COFFEE	
16:00-17:30		Discussion social activity M2M general breakout Lightning Talk Lightning Talk	Break out summary and breakout plan Lightning Talk Lightning Talk Lightning Talk	Social Event/Hiking	Break-out: Social, How (Applications)	
17:30-18:00		Open discussions; Individual Working Groups	Open discussions; Individual Working Groups		Open discussions; Individual Working Groups	
18:00-	DINNER	DINNER	DINNER		DINNER	

Figure 1: Seminar Agenda

also be required to ensure that the M2M devices need not keep track of all the interested subscribers or those might publish data that is of interest to them. Also, introducing cloud-based solutions promises high potential but also challenges regarding reliability, cost and security issues. Use of ICN, SDN and NFV to facilitate these services needs further research and joint activities.

Within the GI-Dagstuhl seminar, the participants identified most relevant, partly overarching aspects to the three main directions above. However, during the discussions most of these questions were addressed to one degree or another. The agenda is depicted in Figure 1. We decided to have the talks mainly in the morning and group work and discussions during the afternoon and evening. All presentation are available online[2].

[2] http://materials.dagstuhl.de/index.php?semnr=16353

Contents

2 Invited Presentations

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

2.1 Context-based support for M2M (Christian Becker)

Christian gave a comprehensive review on research on Context awareness, Context management, Context sources, and Context representation spanning from the very beginning of the research field to recent developments.

After more than two decades of research on the topic, a matured understanding of context modeling, organisation and specific aspects (e.g. data quality) has been developed. Context-computing has been part of major venues such as UbiComp, PerCom, CoMoRea or CONTEXT. Still, the field is swiftly evolving and a number of exciting unsolved research issues exist. Among these are proactivity, deep learning, activity recognition, interoperation between context providers, discovery of context and context providers. In particular, with respect to the research expertise of a group of participants, the question of addressing entities providing context is also of great interest. Are we on the verge from IP to NDN to context-based communication?

2.2 Security for M2M – Block chains (Pekka Nikander)

Block chains is largely a hype term, made popular by the publicity on BitCoin. In essence, the term denotes technologies that allow creation of open, decentralized, undeniable, consistent event logs, which are also often called distributed open ledgers. The good thing about all the BitCoin hype is that it has made the potential of well designed cryptographic protocols well known to a large fraction of the public. However, at the same time the publicity has created a creative frenzy and a set of unreasonable expectations.

In this talk, Pekka introduced the basic technical ideas behind block chains and considered some of the main design choices one has to make when applying block chains to some application area, including identity management, consistency and consensus semantics, and incentives. In the end of the talk, he briefly considered the applicability of block chains to Internet of Things (IoT) applications.

The ensuing discussion largely centered around what is the relationship between block chains and the real world, and what that real world actually is in this context. One of Pekkas main tenets was that block chains without explicit real world connections will eventually die out.

Carsten asked what is the real world? As an example, he stated that for example does Cisco just exist in the digital world, basically implying that the concept of the corporation called Cisco does not exist in the real world, at least if we understand the real as atoms and molecules.

In his answer, Pekka used the concept of money as an example. Money exists both in the real world (in the form of notes and coins, at least), while it is today largely a digital world phenomenon that greatly affects what we are able to do and not to do in the world of atoms and molecules.

Today, money is basically debt that has its roots in the property rights. The debt binds money to the real world.

Considering the history of money, precious metals become valuable since they are relatively easy to carry, measure, stamp, and they have real-life use in the form of jewelry and in tools. Today, gold is so expensive largely due to its relative scarcity and the expenses related to mining new gold. The amount of mined gold is quite small compared to the total gold stored in vaults. BitCoin is based on scarcity that is artificial, created by design. Mining new BitCoins will require more and more energy, and Pekka believes that mining new BitCoins will eventually stop due to this.

From this point of view, the real world is the set of social beliefs that we are sharing. For example, we believe in the government and in the laws and enjoy the consequences, including that we can use money to buy things we desire and may call the police get rid of unwanted people from our property.

A second discussion topic was whether we can use block chains for distributed IoT, without needing to store the whole block chain history in the individual nodes. Pekkas answer was a qualified yes: if the stakeholders of the system are fine with trusting what is already committed into the open-ledger, i.e. what everyone can see in the transactions, then, depending on the design of the information store, one can proceed without needing to store or verify the whole transaction chain. Full history is needed only in the case one wants to verify inputs from very old transactions, and even in that case the full history may be stored in a distributed manner, as long as the nodes collectively have sufficient incentives to keep the full history accessible.

2.3 Services for M2M – DEMIS (Olga Streibel)

DEMIS, the German electronic reporting system for infectious diseases control, is aiming at creating a reliable platform that provides services for epidemiological experts and decision makers in order to support their daily work. The goal is to offer a secure, reliable and fast responding system that can handle the respective data and information in an intelligent way. One of main challenges in DEMIS is the technical conception and implementation of the functionalities required under the constraints given by law and by the stakeholders. Focusing on one of the main components of DEMIS, the terminology service, we tackle the problem of data and information modeling and especially the problem of (data) privacy.

Suggestions from the discussion

Applying additional information sources Whereas the important information is flooding into the system from physicians reports, medical lab data, and the local/communal health authorities, one of additional information sources is to consider taking into account the public available data (i.e. relevant news and reports or users updates from social networks). Thereby the so called trend mining approach (see https://sites.google.com/site/tremitool/trend-mining) can be applicable to some extend. By analysing the public information with regard to the trend(s) contained in it, it might be possible to automatically generate early warnings for the DEMIS system. However, there is a problem with trust in information contained in publicly available sources as e.g. user posts or news updates. Besides that it might be perceived by responsible experts as risky to rely on a trend that is automatically calculated on the basis of publicly available user posts (i.e. Twitter updates). Another possibility could be the consideration of data from external devices, such as measurement devices that being used in medical labs can automatically perform gene sequencing, or consideration of the data from IoT environments. It is definitely a useful and advantageous aspect from the research point of view. From the practical or legacy point of view, it might be not applicable because of legal constraints and data privacy issues. The best is to find a solution for retrieving anonymous environmental data that might be sufficient as, again, an early warning signal for the system.

Handling the privacy and data security

One of the most important issues is the privacy and security of data used in DEMIS. With regard to the three levels model: 1. Serialization, 2. Data modeling, 3. Information modeling (see Carstens slides on the need for rethinking the non-homogen modelling throughout the three levels on which we handle data) there is a need for clarifying at which level and to which extent the privacy is important. DEMISs terminology service relies on sophisticated ontology models and scheme mapping all together called the intelligent and adaptive knowledge base. While creating the respective knowledge base we should consider the relevance of the privacy in the model. Furthermore it is important to ensure data security while retrieving and delivering DEMISs data (i.e. in planned LinkedData endpoints).

Discussion summary

Public health data management grows in importance nowadays. While creating systems that handle the public health data we always tackle the problem of data privacy and security. Being able to use the data streams from IoT might enhance our ability to offer sophisticated analysis and reliable results, however there is a need for finding proper solution for data privacy. Data privacy (see the discussion in the privacy breakout session) issue needs to be considered under the existing law (e.g. the German law in this case). However, even while having a strict regulation about using the sensible data (e.g. patients data) there should be possibility of including the information built on these data into the system without violating the respective privacy (therefore there is a need for defining and considering different aspects of privacy in respective context). Moreover, it is important also how the information is presented to the respective decision makers (adaptive and flexibel UIs) and how much of the information can be gathered automatically. We should try to get away from systems where physicians type in the information/reports manually. The goal could be to create adaptive and self-thinking systems that reliably support physicians and experts in the public health. (Disclaimer: the thoughts above are the results of the discussion with regard to the relevant research problems and aspects in the DEMIS system.They are not the official DEMISs project point of view.)

3 Parallel Group Work

We identified relevant topics with all participants and worked on these in topical breakout groups.

3.1 Definition of M2M

In this session we tried to derive a useful definition for machine-to-machine (M2M) communication. We focused on a definition that serves the needs of the participants instead of coming up with a very generic statement.

We identified the following indicator as important to characterize M2M communication: At least one endpoint should have a relationship to a physical object. As an alternative view both endpoints need to be machines. However, we agreed that a scenario in which both endpoints are not a physical end point will not meet M2M communication. Regarding the type of devices there are no restrictions. In particular, tiny devices may or may not be involved on the physical side.Application scenario may imply lower software upgrade rate.

After discussing indicators we analyzed potential implications that arise for M2M communication. Most importantly is interoperability. We also agreed that where data are interpreted by a machine (as opposed to a human), there may need to be stricter semantics attached to the data. We further agreed that the decision process is different. Decisions may need to be framed into *policies* so they can be taken without the humans present. However, this may become more open with artificial intelligence present.

We did not conclude on the questions, if M2M communication leads to more predictable communication patterns. As (i) machines serve a specific purpose and (ii) humans are not directly involved, one might argue that the event space is more limited (i.e., less random). However, when sensing the environment human behavior may still lead to unpredictable conditions, for example.

Finally, we question the differences between IoT vs. sensor/actuator networks vs. M2M vs. cyber physical systems (CPS). A CPS does not need to be (intra/inter-)networked—but it often is. The term Internet of Things (IoT) might be used to imply that communication is across multiple autonomous systems, i.e., inter-domain

3.2 RIOT (tutorial and discussion)

The breakout session consisted of two parts: a brief introduction into RIOT and a hands-on tutorial on RIOT

RIOT is an open-source operating system for the Internet of things. It targets low-power, memory constrained devices connected over low-power and lossy networks. It is based on a microkernel architecture with an energy-efficient and real-time capable tickless scheduler and a multi-threading programming model. Two of RIOTs key design principles are modularity and interoperability through open standards such as POSIX or IETF protocols. RIOT is currently developed by a world-wide and very lively community of more than 100 contributors.

The tutorial started with an explanation about the minimum requirements to develop an application for RIOT and an overview over development, experimentation, and deployment facilities. It was shown how to install and interact with RIOT on real hardware (an Atmel SAMR21-Xpro, based on an ARM Cortex-M0, was provided as an example to the participants), how to use RIOT in an emulation environment using the native port, and accessing the IoT-LAB testbed in order to deploy and execute experiments.

The first tasks in this tutorial explained how to develop a simple application for RIOT, using the shell and implementing simple commands. After an introduction into the threading, IPC, and timer systems, the participants learned how to use RIOTs networking capabilities based on the GNRC network stack. The participants wrote a basic UDP server-client application to exchange small text messages between the local nodes and inside the IoT-LAB testbed[3]

3.3 Sensing

In this break-out-session, we first tried to define the data we are talking about. In particular, the data is often in the form of time series from which some might be latency sensitive and some might require time-synchronisation. It can also feature different accuracy (e.g. event based vs. full accuracy), price of the sensor, resolution, and metadata (placement, type, . . .).

It is also possible that sensor data from multiple sources are fused to a single time series. This type in which the data is in, has implications on communication and encoding of the data. Another way of seeing this is to consider services to provide data from sensors which is either fused or from individual physical devices. For several reasons it might be desired to pre-process and maybe reduce the data forwarded or shared from a sensor. Important aspects are

energy computation is usually cheaper than communication

privacy reveal less information (but see 'anonymization' below)

Indeed, standardized sensor data formats (a la OGC SensorML, IETF SenML, . . .) exist. It is desirable to exploit common aspects with these in advance of defining a novel description for a new kind of sensor. Privacy is an important issue in the sharing and pre-processing of M2M sensor data that is addressed

[3] Material and requirements for the tutorial: http://tutorial.riot-os.org/README.md For testbed access you need to register at https://www.iot-lab.info/testbed/signup.php and install the cli-tools as described here: https://github.com/iot-lab/iot-lab/wiki/CLI-Tools-Installation

by several recent projects [4] [5] [6]. Also, noise is an inherent property of sensor data from M2M sensors which might also be exploited intentionally in order to protect individual privacy. One problem with such approaches is that it is potentially always possible to deanonymize given enough data and over sufficient time. This discussion also sparked ideas that have further been discussed in the breakout session on privacy. In particular, it is a valid question if an unrestricted access to all data by everyone (as opposed to only individual authorities) would actually resolve the privacy issue with sharing data.

Example: Automatic generation of a Shared Secret

We discussed an example application in which on-body devices are paired conditioned on a fingerprint from acceleration sequences that capture human gait. The problem experienced was that acceleration sequences significantly differ over different body parts which makes it difficult to retrieve a high-entropy key from the gait sequences. A suggestion by the participants was to rephrase the problem zero-knowledge proof help. In this way, sensors would together arrive at the same secret.

Also, participants pointed out that already paired devices could cooperate in the pairing process and whether it would be possible to extract privacy-related features such as person height from the gait. Separating meta-data from measurement data may help to improve (weak) privacy.

In general, the authentication of devices relies nowadays usually on laws (which limits flexibility on the devices) or silo solutions. For instance, the company Apple has an infrastructure for device authentication which is based on a single trusted party (apple).

Opportunities and outcome

It is important that we become able to plug things together in flexible ways by making sensor interfaces interchangeable. Also, we should understand the range of requirements since this makes it easier to design systems that can talk to each other as required. So, different use cases should be profiled, e.g., step counter but also a taxonomy is necessary to classify different sensors and sensor classes. Incentives for vendors to comply to these requirements might be imposed by regulation but also opportunity to participate in a particular market may be one. Also, the output format needs standardisation (e.g. certainty, confidence, data quality) since this makes it easier to collect corpuses of data to use for others.

3.4 Usable security

In this break-out-session, we talked about useful tools and methods for usable security. The group has been composed from experts in various fields in the

[4] http://www.databoxproject.uk/?page_id=31

[5] http://hubofallthings.com/

[6] Remove privacy-invading camera images from lifeloggers: http://private.soic.indiana.edu/projects/cameras/

domain, who then gave brief topical introductions to seleted issues in usable security and discussed the implications and opportunities with relation to M2M. A general introduction to the security issues faced in M2M environments is provided in [31].

Fuzzy encryption – ad-hoc audio-based device pairing (Stephan Sigg)

Fuzzy cryptography, such as, for instance, fuzzy commitment schemes enable the generation of shared secrets from noisy or only partly identical input. This approach can be used for pairing in M2M scenarios, conditioned on data sensed in the same context, such as audio or RF in proximity, or, for instance, acceleration sequences for devices moved jointly [28, 30]. A good introduction to the topic is provided in [32] as well as [18, 17, 16, 10, 11, 26] for different realisations of the general concept. If the sequences extracted on two devices are close enough to one another, devices can can obtain the same secret key. Otherwise, they can not. The 'threshold' of error-correcting code schemes can be configured to correct the difference between binary sequences.

It has been mentioned that one alternative is to use 'commit and exchange of hash'.

In any case, the entropy of the data utilised need to be sufficiently high. Also, work how to attack those schemes has been presented, for instance in [27, 15]

Actor model for Web of Things (Carsten Bormann)

The actor model has been presented for authenticated authorisation of M2M devices. In particular, it provides a solution to providing strong authentication routines also for restricted devices by separating authentication and authorisation into two levels. It was discussed that this model nicely matches to the above concept of fuzzy cryptography, where additionally, proximity can be conditioned for authentication to work. Also, it was mentioned that in this model still the issue with revoking the authorization remains.

Attribute Based Encryption (Börje Ohlman)

Attribute based encryption enables that objects to be published are secured already at the sensor where they are produced and not along the gateway. Attributes can be for instance roles (teacher, student, ...). The approach utilized public key encryption end differences between a key policy and ciphertext-policy, dependent on whether the access policy is defined in the cyphertext or in the keys. For M2M purpose especially ciphertext-policy is relevant. It enables to encrypt once and allows access for an unlimited amount of recipients conditioned on e.g. membership to company. In implementation for this might use symmetric encryption on the sensor but give those with the right attributes access to the symmetric key via attribute-based encryption. A 128 bit security level has been proposed. The scheme boasts the benefit that it does not require the need to keep the data in a safe place one produced but instead secures the object itself which is then stored all over the place and access is given via attribute-based encryption so that some people get access but others not. Another benefit in the ICN case is that one object looks the same in the cache and not different for different encrypted versions of the same object, which saves storage space.

Forward security (Dominik Schuermann)

An issue with secret keys is that, once the key is stolen by an adversary, this will usually disclose all past and future data to the observary. This generates an incentive not to use secret keys that last for a very long time. Changing the key frequently hinders the adversary to read what has been communicated in the past (Forward secrecy). A suggestion is to change the key with any message exchange. In this case, even if someone steals the key inbetween, provided that authenticated Diffie Hellman was used, the attacker can not follow up with the changes in keys since Man in the Middle attacks are not possible. However, problems to decrypt messages could arise when ordering of messages is confused as it is frequent in current communication protocols (e.g. TPC). A solution to this problem is the use of Hash-Ratchets. This is a deterministic one-way function to derive the next key. This concept still provides forward secrecy, but once stolen, backward secrecy is not guaranteed With the use of double ratchets, also backward secrecy can be established. This means that, in addition to using a function to deterministically derive the next key, the protocol starts over with a new session key for each new session.

User study on perceived complexity (Dominik Schuermann)

We discussed the results presented in a recent USENIX security work [7]. In particular, it was observed that approaches to manually compare hash sequences utilise different encodings of the hash and that hexadecimal encoding appears to be fast for human beings to grasp and with low error rate. It turned out that numerical sequences are best while the often used hexadecimal encoding is the worst among the considered encodings. In addition, a new representation of password hashes in the form of semantically correct sentences has been proposed and tested with good results.

3.5 ICN

We started with a discussion of ICN in general with those who are not working on ICN. In particular, we discussed the abstraction in ICN and what belongs to the application layer. Two application scenarios have been discussed: Vehicular networks (sensors and cloud) as well as Industrial safety. Aspects considered were, how useful ICN is for other application scenarios, especially service oriented platforms and how we can demonstrate that ICN is useful for 'non large scale content oriented' networks. It has been mentioned, that active networks, and pub/sub were there, but largely unnoticed. However, with larger corporate/controlled networks, ICN paradigms could be implemented. The question that remains is how the applications look or would be built and what paradigms will be required.

3.6 Privacy and some related topics

A breakout session on privacy and related issues was held on Thursday after lunch. As privacy is such a large area and hard to focus on in a brief time, a goal was to create an idea of the most pressing topics people are working on, find common themes, and try to bring up with some fresh thoughts related to the common themes.

Using one of Finnish philosopher Esa Saarinen's teaching methods (see [20] for some reflections), we collected thoughts from all participants, in trying to gain common shared ground. Briefly stated, the method starts with each participant individually reflecting a topic, then sharing the thoughts with someone in the group that he or she does not know (that well), and then gradually sharing the thoughts in larger and larger groups. The method is part of what is called ”The miracle of Lecture Hall A” in [20].

In the first round, two topics were raised above others:

- Correlation of independent data streams, especially in the context of location tracking and using of stable identifiers in communications

- Impact of information availability

As is well known, correlation of independent data streams allows one to associate various identifiers with each, often with quite high accuracy, especially if the data streams a long. Consequently, when correlated, any designs where only the identifiers in a data stream are anonymous, while other data is presented openly, is likely to reveal quite a lot about the identifier holders (see e.g. [3]). The other topic, impact of information availability, had already been discussed briefly in various occasions. As a working hypothesis, it looks like that the value of information can be divided into two bins:

- Direct utility value, or what new methods or approaches the information allows

- Value derived from information asymmetry, or some parties having information others don't have

It is noteworthy than when information is made public, its overall utility value increases, as more people can utilize the information, but at the same time the value (rents) derived from information asymmetry shrink.

What comes to information asymmetry, in our current society there is a large information asymmetry favouring companies over individuals and a (perhaps smaller) asymmetry favouring rich over the poor. For example, [19] claim that 'information asymmetry is the very foundation on which the existence of elites is built and possibilities of strengthening that asymmetry will be enthusiastically sought'.

In the ensuing discussion round the group converged to the concept of property rights, comparing existing property rights (and associated social conventions) related to real life property, such as houses and bicycles, vs. the largely-absent property rights regarding personal data [24, 29]

More generally, it was clear that proper privacy management needs *both* technical solutions *and* properly enforced regulation.

On a second round of using Saarinen's method, a number of aspects were highlighted:

- Opting in and opting out, with the impact of defaults
- Information asymmetries in current data collection practises
- Ownership of personal data stemming from relationships
- Inalienable property rights with respect to data

Firstly, for any personal data, there should be a probably inalienable right for choosing whether to opt-in or opt-out in any collection of such data. It is notable that while such regulation exists under many jurisdictions, in practise the regulation is poorly enforced, and the regulation is often insufficient. Furthermore, the defaults are here very important, as most people don't bother to change their defaults [14].

Secondly, a large part of the current problems are related that people are not aware (and often cannot be aware) of what data is being collected about them. While in some jurisdictions there are clear regulations about explicit personal databases, these do not usually apply to less-explicit personal data collection nor to other jurisdictions. This is also partly a technical problem; for example, the browsers could do a much better job in how they represent cookies to the users [12, 22]

Thirdly, there are often cases where personal data is born as a result of more than one person acting together. From a data property rights point of view, it is not immediately clear 'who' should 'own' such data.

Finally, on the inalienable rights the underlying thought was mostly that we are no longer allowed to sell ourselves (or our family) as slaves, as was somewhat common a couple of thousands years ago. In the same way, we shouldn't be allowed to give up some rights related to our personal data [13, 24, 4].

Towards the end the discussion diverted to owner's rights vs. owners responsibilities, especially with respect to limited liability companies (LLC). In essence, an LLC limits the responsibility of the owners of a company, in practise even in the case the company commits to crimes. This creates the perverse incentive of a company breaking the law surreptitiously, thereby gaining profits and handling out (some of) the profits to the shareholders. According to today's laws and practises, the owners are not punished for this anyway, usually not even through recollecting the illegally gained profits back. (Consider e.g. [23])

As a though experiment, we considered a change of the juridical practices where the shareholders of a company are punished (e.g. heavily fined) for any crimes the company commits to (even after the company may have been liquidated through a bankruptcy), in addition to punishing the employees who took part in the crimes. This would create a stronger-than-today incentive for the shareholders to make sure that the board and management act according to the laws. This should be considered especially in the context of shareholder primacy[7].

3.7 On the 'How'

In this breakout session, we discussed the "Haystack" application by the Berkeley team that could be installed on smartphones and used to analyze if any application is violating privacy. We also discussed if key revocation could be performed and what issues arise. Some scenarios where key revocation could be useful are: a) CDNs/Content providers who want to revoke access rights to those who unsubscribe; and b) in case of employees who leave the organization (e.g. police men). On the application front, we agreed that it is important to start by designing application specific stacks such as network stacks. As we design more and more application specific stacks, we will be able to identify abstractions/commonality/generic-aspects that could then help in designing a generic solution. Some application based ICN designs that were mentioned were [2, 5].

3.8 Advice for PhD students (Carsten, Christian, Jörg)

Targeting especially the younger participants, senior attendees provided advise on how to conduct PhD studies and where to set the emphasis. In particular, questions discussed where

- How to do research (think in small steps and start from abstract perspective)
- Marketing of the own story (be able to tell the main points in few sentences)

4 Lightning talks

Each participant has been asked to prepare a lightning talk on their research topic. These talks have been presented based on interest of participants.

4.1 Distributed Infrastructure for IoT (Carsten Bormann)

Carsten Bormann gave a lightning talk about *Distributed Infrastructure for Internet of Things*. He started with clarifying why a distributed infrastructure is important. In addition to the observation that it is easier to operate, a distributed infrastructure is helpful to avoid rent-seeking opportunities. A major challenge for such systems is the scalable registration of protocols. An IANA

[7] https://en.wikipedia.org/wiki/Shareholder_primacy

model for the M2M ecosystem will very likely not work. Alternative approaches might involve blockchain implementations such as Bitcoin.

During the discussion the question come up how we can store a blockchain on constrained devices, which only have very limited memory resources. An approach to solve this would be to store the data on a sufficiently equipped device that manages other IoT devices (e.g., border router). This requires a trusted connection from the device to the IoT devices, which exists in many application scenarios as devices are owned by a person.

It is worth noting that even when blockchain is not synced fully to every device, it is still distributed, but not autonomous with a full mesh.

We also discussed mobile code. Usually, mobile code is deployed to increase availability and to reduce latency.

4.2 Authentication based on egocentric videos (Le Nguyen)

We aim to generate image-based passwords, which are temporary and personalized, from users contexts. More specifically, a small camera attached on the human body captures the users activities and locations from the first-person point-of-view. The videos are then processed using segmentation and clustering techniques. In our current research, an authentication password contains four images from distinct video segments. The user is required to re-arrange those images in the right chronological order to log-in to the personal devices (such as smartphone or tablet).

Due to the origin of the images, our passwords are dynamic and adapting according to the variation in users activities. The proposed algorithm can capture various level of changing, ranging from applications on computer screens to moving between different locations. Thus, it is challenging for an attacker to break the passwords if the user is not followed suspiciously everytime. Research questions for future improvement: a) How to protect the privacy of the user when the images are shared via some server; b) Where to process the video (edge or on device); c) How to securely transmit between devices; d) How to authenticate the device to pair to in the first place; e) What happens for routine cases (going from home to work everyday) or in the case of a conference where more than one participant have the same routine. During the seminar, a user study has been conducted: User study link: http://ambientintelligence.aalto.fi/passframe/UserStudy/

4.3 Emotion sensing using radio signals (Muneeba Raja)

Emotion sensing has a well known research topic in field of medicine, computer science, machine learning and human-computer interaction. So far, emotions are detected using facial expressions, speech, text or physiological signals. All these modalities require extreme involvement of humans. Sources such as video cameras, wearable devices and social media data is being used for human emotion mapping. We propose a technique which leverages radio frequency signals for emotion sensing [25]. The main idea is to track human body movements and gestures which depict certain emotions. For example, the body gestures of humans in state of fear or anger are different from the normal state. And they are indicators of strong emotions.

Application in car environment. We have proved our idea by performing a practical experiment of differentiating between angry driving behavior and neutral driving behavior based of body movements.

In relation to machine-to-machine communication, the application follows pubsub model. With RF activity recognition we create a new type of sensor node. This sensor node will be able to detect human gestures/emotions/behavioral states. And it can publish this information and other devices can use this information. So, for instance, the emotion recognition module in the car detects anger state of its driver, and publishes this event. The smart home will receive this event and do the required action, e.g., play some soft music and warm up the home even before the person has arrived home.

Issues to address:

- Accurate gesture recognition
- Security / privacy (tracking human emotions)
- Data handling
- Other approaches than RF
- Is it suitable for vehicles

4.4 Private Proximity-based Services (Michael Haus)

First of all, we introduce the concept of Proximity-based Services (PBS) and which is the important difference to Location-based Services (LBS). PBS use additionally the relative distance between two or more users or between an user and an object. The goal is the automatic identification of physical proximity to points of interest (POIs). Thus, we are able to share information within time-limited semantically grouped people. Many solutions focus on location information model proximity. We aim to utilize other useful information, such as audio and vision to infer proximity between POIs. For example, an use case to enable automatic calendar sharing between people attending the same meeting. Besides that, there are also industry efforts in the PBS domain to standardize D2D proximity services for LTE by the 3GPP group or the Wi-Fi Aware Alliance. We received the Google Internet of Things (IoT) Technology Research Award Pilot and based on these devices we are currently building a testbed to establish useful proximity services. The user privacy is another crucial aspect for PBS which are based on sensor data from mobile devices. For instance, a proximity based application uses audio data to infer users in the vicinity. The sound data is very sensitive regarding user privacy. Therefore, we have to secure the local sensor data.

One way is to use a Personal Data Store (PDS) to improve the users privacy by controlling the release of sensitive information. These systems allow a fine-grained data access control, such as how the sensor data is accessible and at which frequency, e.g. location updates. Another aspect is to adjust the sensor data before releasing them to a potential attacker. The outside world has no access to the raw sensor data, the PDS only provide processed results based on the raw data set. Projects in that area are the Databox Project from Queen Mary University and the openPDS system from MIT.

The second possibility to secure the user data is to hide the data from potential adversaries. Thereby, the data itself is not adjusted as by a PDS, the content is protected by encryption. Private Proximity Testing (PPT) is often considered as an instance of a secure multi-party computation (SMC) problem. The PPT problem is reduced to private matching problem (i.e., private equality testing (PET) or private set intersection (PSI)). In this problem, each party holds a set of inputs and needs to jointly calculate the intersection of the input sets without revealing any further information. The two main techniques for PPT are homomorphic encryption to directly perform computations on ciphertext (e.g. location information) and garbled circuits, in which one party prepares an encrypted circuit computing function; the second party computes the output of the circuit (e.g. proximity) without learning any intermediate values.

4.5 SoftOffload and Securebox (Aaron Yi Ding)

Two initiatives are illustrated at the seminar aiming to make the network edge smarter and safer, through data traffic management and IoT security. First, SoftOffload [9] is an open-source software defined platform for achieving intelligent mobile data offloading. The platform consists of an extensible central controller, programmable offloading agents, and offloading extensions on mobile devices. Motivated by the measurements of energy consumption on smartphones, we propose an effective energy-aware offloading algorithm derived from MADNet [8] and integrate it to SoftOffload. By enabling collaboration between wireless networks and mobile users, SoftOffload can make optimal offloading decisions that improve the offloading efficiency for network operators and achieve energy saving for mobile users. To enhance deployability, we have released our platform under open-source licenses on GitHub.

Second, Securebox is an affordable and deployable platform for securing and managing IoT networks. The design is motivated by the observations that IoT security is becoming more challenging due to the device limitation, interaction with physical space, and cross-device dependency. To make the problem even worse, end users typically lack the expertise and budget to manage the IoT devices deployed at home/office environments. In this regard, Securebox targets an alarming spot in the fast growing IoT industry where security is often overlooked due to device limitation, budget constraint, and development deadline. In contrast to existing host-centric/hardware-coupled solutions, Securebox empowers a cloud-assisted charge for network service" model dedicated to budget and resource constrained IoT environments. Based on its cloud-driven and modular design, Securebox allows us to flexibly offload and onload security and management functions to the cloud and network edge components. It offers advanced security and management services to end users in an a ordable and on-demand manner. Furthermore, Securebox can ease the upgrade and deployment of new services to guard against abrupt security breakouts. To demonstrate Securebox, we have implemented the platform consisting of a plug-n-play frontend, a Kubernetes-powered backend cluster, and a smartphone mobile application. Based on the testbed evaluation, we show that Securebox is robust and responsive. Its collaborative and extensible architecture enforces rapid update cycles and can scale with the growing diversity of IoT devices.

To sum up, the network edge has been expanding exponentially in terms of device diversity, scale, and traffic volume. On the other hand, those challenges have generated good research directions to rethink the protocol and system design for building a more robust and scalable Internet ecosystem.

4.6 Proactive caching in ICN-VANETS (Dennis Grewe)

The Bosch portfolio offers products from different domains such as automotive, smart home, household appliances and so on. In 2015 Bosch introduces the Bosch IoT cloud to connect devices from all of the Bosch domains together. By connecting a lot of devices with an infrastructure, the company questions if the current host-centric, end-to-end communication model of todays Internet is suitable for communication participants interested in same data (e.g. weather forecast within a geolocation aggregated by sensors within this area).

So Bosch started an activity looking into the Information-centric networking paradigm and its concepts which provides natural support of mobility as well as in-network caching. First, use cases from the vehicular networking domain are used to evaluate ICN concepts. One of the uses cases is introduced by consuming data from a cloud infrastructure periodically e.g. to optimize the battery consumption of an electronic vehicle. When using the standard ICN caching approaches, it can be seen that a mobile consumer may have lost the connection to the initial Access Point (AP) and hence is not able to receive the data. As a result, it has to repeat the request at the next entry point. In worst cases, the response does not reach the mobile node and hence is being repeated over several time (we call it mobile node delivery problem). Such problem decreases the resolve rate and thus increases the overall bandwidth used. One of the research questions are: How to cache data proactively at the AP to increase the resolve rate while keeping the average number of cached copies within the network to a comparable level? How to achieve such mechanism without changing the communication stack?

Furthermore, there are some Bosch activities trying to overcome technological silos by using Eclipse Vorto (www.eclipse.org/vorto/). The toolkit allows to describe the attributes and the capabilities of real world devices such as sensors or actuators as part of so-called Information Models. Such models are used by Code Generators to integrate devices into different platforms. Based on this toolkit, we introduced a harmonized and generic cross-platform interface for vehicles and the cloud using Eclipse Vorto [33].

4.7 Scale IoT computation on the Edge (Nitinder Mohan)

Internet of Things typically involves a significant number of smart sensors sensing information from the environment and sharing it to a cloud service for processing. A recent study by National Cable & Telecommunications Association (NCTA) assumes that close to 50.1 billion IoT devices will be connected to the Internet by 2020. Even though the traditional cloud model is able to handle processing requirements of IoT data, transporting data over large geographical distances imposes a significant network delay which further impacts the overall processing time.

Various architectural abstractions, such as Fog and Edge computing, have been proposed to localize some of the processing near the sensors and away from the central cloud servers. The objective of Fog cloud is to pre-compute the data while routing it to the central cloud such that the Fog resources perform low-latency computation/aggregation in the network and leave the heavy computation to central cloud. On the other hand, the Edge cloud proposes a consolidation of human-operated, voluntary resources such as desktop PCs, tablets, smart phones, nano data centers as a cloud which lie at one-hop distance to sensor nodes.

In this talk, I gave an overview of the Edge-Fog Cloud which distributes task processing on the participating cloud resources in the network. As the name suggests, the outermost layer of the Edge-Fog cloud is composed of a large number of volunteer, human-operated edge devices connected via ad-hoc network chains. The inner layer is composed of a dense network of Fog devices with high compute capabilities.

I further gave an overview of Least Processing Cost First (LPCF) method for assigning the processing tasks to nodes which provide the optimal processing time and near optimal networking costs. Our solution is based on the observation that by first minimizing the processing time, we can achieve near-optimal networking costs in polynomial as opposed to exponential time complexity. We evaluate LPCF in a variety of scenarios and demonstrate its effectiveness in finding the processing task assignments.

4.8 ICN for IoT (Sripriya Adhatarao)

The speaker presented a talk on the possibility to use Information Centric Networking for IoT. The IoT traffic pattern is mainly query/response and scheduled updates. There is basically a single entity like Base Station (BS)/Base Station Controller (BSC) that collects the information from the sensing nodes/devices. Based on these observations, it was suggested that ICN is a more suitable candidate for deploying IoT. Since the users in the IoT domain are also interested in content irrespective of who provides them or where it comes from. One additional and crucial requirement for IoT devices is Security and it can be easily achieved with ICN since the Content is signed by the publisher. ICN can significantly improve the efficiency of the IoT devices along with reduction in energy consumption. Naming is another critical aspect [1] that needs to be resolved and the speaker pointed to the need for different naming schemes in different parts of the network.

An important concern was raised by the speaker that existing ICN architectures are too heavy for the IoT devices. Lighter versions of the proposed ICN architectures like CCN-lite could be used for IoT. A question was raised regarding the benefit of running CCN-light on top of wsn? which needs further analysis. Further, the suggestions were made in the proposal about a need for a mechanisms to integrate Sensor Networks to the Internet. Gateway is a good direction to think about integrating the Sensor networks to Internet. We need to clearly define the functionalities expected to be supported in the Gateway. However the discussion of gateway capabilities exceeds the scope of this work. The speaker was suggested to look at the 6LoWPAN border routers since they have well defined Gateway functionalities that can be referred. The proposal also spoke about naming schemas for IoT networks. Suggestions were made about the possibility to verify the Header Compression techniques. But, we have to remember the limited computation capacity and overhead on the resource constrained devices.

The proposal also mentioned that Pub/Sub [6, 5] is a better suited communication model for IoT networks. It was also pointed out during the discussion that it is important to differentiate between the IoT and the Sensor networks that are expected to be connected to the Internet. One of the suggestion was to refer to the less constrained network as either Sensor Networks or lossy/low power networks. The group also agreed on the need to compress the full fledged ICN protocol features to only those that are important and needed in the IoT networks. It was also decided that we have to design an architecture where the deployment and configuration of Gateway should not be too stringent to the applications. As it may lead to frequent re-configuration of the Gateway. The speaker was also suggested to look at the implementation of NDN using RIOT for evaluating the proposal.

4.9 Smartpone data sharing and processing (Ioannis Psaras)

The talk presents a framework for sharing processing and storage resources on mobile devices. The authors experimented with WiFi direct for device communication which worked more robustly than Bluetooth. Shared apps should not involve personalized content, such as routefinder apps where people have no concerns sharing it. However, there is an issue when people decide which applications they do and dont advertise, since this way fingerprints are created for each user by which the person might be identified. A possible solution of this might be to restrict the sharing to friends.

Since such kind of ideas have been there for a while, it is important to stress the new aspect, such as device interaction. Also, the concept to involve processed information and/or computation appears to be novel. Related to this might be the discussion among Hadoop people who regarded named information but did not consider much implementation details. It should be investigated though, if this approach is actually cheaper and more efficient than finding the app online and installing it locally for use.

4.10 Gait-pairing for Wearables (Dominik Schürmann)

In this work, we introduce a way how to pair devices worn on the same body. As a physically separating property of the human body we utilize gait sequences recorded using accelerometers. Our early results indicate that it works with relatively little pre-processing for sensors worn on the waist. Fingerprints generated from sensors attached to the left and right waist, extracted from a gait data set of the Osaka University, were sufficiently similar to each other for usage for authentication. For this quantization step a mean gait cycles has been calculated independently on each device for the recorded time span. Then, 4 bits are extracted from every recorded gait cycle by comparing it against the mean gait cycle. If the energy is above the mean gait cycle a 1 is extracted, else a 0. While this works using the Osaka data set, we still need to figure out how to better re-orientate sensors worn on more dynamic body parts, such as forearms. For this evaluation, we are using a data set from the university of Mannheim, where sensors have been attached to different body parts.

Looking at different gait cycles, it can be observed if a person has medical conditions, e.g., if the person hinges. Also, the gait cycles changes drastically if a person is getting injured, e.g., breaks a leg. While the protocol no longer works when the person is not able to work, small injuries have no effect on the success of the protocol, because our protocol only uses the gait cycles for a specific time period for pairing at this point in time. In contrast to authentication methods for unlocking phones, we do not rely on historical gait cycles, i.e. no enrollment process.

Besides the presented related work, we also looked into the algorithms for step trackers. They use similar but much more coarse methods of processing the sensor readings. In contrast, we try to extract the unique features in between steps, they are only interested in the number of steps.

In our threat model, we discussed how the pairing can be attacked, especially how a persons gait can be forged. The easiest way to attack the protocol is by attaching a malicious sensor to the victims body, maybe by hiding it in his/her clothes (jacket). Another way would be to record a video of the persons gait cycles at the moment of pairing and then reconstruct the accelerometer readings from that. The protocol can not be attacked by using gait cycles recorded at a different time and context. As discussed before, no historical data is used for pairing in contrast to authentication methods, such as phone unlocking.

4.11 Attribute Based Encryption (Börje Ohlman)

Attribute-Based Encryption (ABE) is considered to one of the most promising ways to be enforce access control in Information-Centric Networking (ICN). As ICN is well suited for the Internet of Things (IoT) the question of compatibility between IoT and ABE arises. In IoT there is the resource constrained devices and in ABE there is the computationally expensive operations. This presentation discusses two ways of how CP-ABE can be applied in systems with resource constrained devices. The feasibility of the two systems are evaluated based on experimental results of ABE on a resource constrained sensor. The most suitable solution is determined by the computational power of the sensor, the maximum policy length and the time requirements of the sensor application.

Some example features of Attribute-based encryption are: a) attributes such as roles (teacher, student); b) Public key encryption; c) Key Policy and Ciphertext-policy. The access policies will be present in the keys. The idea is to use symmetric encryption on the sensor but give those with the right attributes access to the symmetric key via attribute-based encryption the possibility to have 128 bit security level. This facilitates to secure the object already at the sensor where it is produced and not along the gateway. The advantage is that encryption needs to be done once and unlimited amount of recipients (conditioned on e.g. membership to company) could receive it. This implies that there is no need to keep the data in a safe place but instead secure the object itself which is then stored all over the place and access is given via attribute-based encryption so that some people get access but others not. Another benefit in the ICN case is that the storage looks the same for one object and not different for different encrypted versions of the same object

It is possible to employ ABE on sensors if the policies are limited in size and if time is not major concern. The time issue can be potentially circumvented by using the same session key for a certain amount of time and refreshing it periodically. This would reduce the time to the same as using ordinary symmetric encryption with the addition of the cost of periodically performing ABE encryption. The limiting factor of the feasibility of performing ABE on resource constrained devices is the RAM size of the device. The gain of performing the ABE operations on the sensors is largest in a multiple authority scenario as this removes the single trusted third party. In a single authority scenario the authority ABE system is preferable because it will be faster, supports arbitrarily large policies and there already exist a trusted third party. However it requires end-to-end communication between sensor node and authority.

4.12 End-to-End Authentication for IoT (Thomas Schmidt)

Authentication of smart objects is a major challenge for the Internet of Things (IoT), and has been left open in DTLS. Leveraging locally managed IPv6 addresses with identity-based cryptography (IBC), we present an efficient end-to-end authentication that (a) assigns a robust and deployment-friendly federation scheme to gateways of IoT subnetworks, and (b) has been evaluated with a modern twisted Edwards elliptic curve cryptography (ECC). Our early results demonstrate feasibility and promise efficiency after ongoing optimisations. Refer to [21] for more details.

The main takeaways of the talk were: a) End-to-end authentication for constrained IoT devices; b) Independent local authorities (possibly offline); c) Federated trust: bootstrapped by crypto-based identifiers; d) Revocation of trust triggers local renumbering; e) Lightweight implementation based on twisted

Edwards Curve (25519), RELIC and RIOT. Furthermore during the discussions, the author mentioned that their work could support application use-cases such as authentication end-to-end (E.g., fake origins can be prevented in routing); and to protect other infrastructure-based services like application directories, dispatcher and etc. The presenter also mentioned that there is one TA per subnet. A research direction that requires further exploration is the application semantics that can be possibly derived from this topology-centric authentication mechanism.

4.13 Neutrality (Pengyuan Zhou)

Network neutrality has become one of the hottest topic lately. After FCC promoted it, lots of enterprise and researchers are participating in the discussion about how to realize it. Facebook built Internet.org to help people all over the world to access internet freely. Yet, lots of people think its against neutrality because it uses network resource with higher priority, which is not fair to other service providers. There are lots of open issues related such as: a) How to realize network neutrality; b) How to motivate the service providers take part in the procedure, since their profit may decrease; c) How ICN could help neutrality; and d) How to realize neutrality in IoT.

4.14 Data formats (Carsten Bormann)

The presenter gave a short overview of data formats and some insights into how these data formats could be represented by JSON and other formats. A problem that was highlighted is that there are too many different data formats in the IoT to translate between and that different interaction models are required. A suggestion by the audience was to use machine learning model based training to perform the transformation

5 Conclusions and next steps

Participants featured a mix of senior and junior researchers from both academia and industry, as well as standardization bodies, participated in fruitful dialogue. The possibility to organize a follow-up meeting after 18 to 24 months has been raised and discussed. In particular, a good opportunity would be to organize a next meeting in the frame of a Shonan Seminar in Japan, since numerous project collaborations between the participants and Japanese researchers exist. The organizers will investigate this opportunity and consider to submit an application to the National Institute of Informatics. The organizing team also received valuable feedback.

6 Participants

- Sripriya Adhatarao, Univeristy of Goettingen

- Mayutan Arumaithurai, University of Goettingen
- Christian Becker, University of Mannheim
- Carsten Bormann, University of Bremen
- Aaron Yi Ding, TU Munich
- Dennis Grewe, Bosch
- Oliver Hahm, Inria
- Michael Haus, TU Munich
- Felix Juraschek, MSA Auer GmbH
- Dirk Kutscher, NEC, Heidelberg
- Nitinder Mohan, University of Helsinki
- Le Ngu Nguyen, Aalto University
- Pekka Nikander, PulseOn
- Joerg Nolte, Brandenburg Technical University
- Borje Ohlman, Ericsson, Sweden
- Ioannis Psaras, UCL, London
- Muneeba Raja, Aalto University
- Thomas Schmidt, HAW Hamburg
- Dominik Schuermann, TU Braunschweig
- Stephan Sigg, Aalto University
- Olga Streibel, Robert-Koch Institut
- Matthias Waehlisch, FU Berlin
- Xiaoyan Wang, National Institute of Informatics
- Lars Wolf, TU Braunschweig
- Lei Zhong, National Institute of Informatics
- Pengyuan Zhou, University of Helsinki

Acknowledgements. The organizers would like to thank the participants for their contributions. Special thanks to Olga Streibel for her help during the reviewing of the manuscript.

References

[1] Sripriya Srikant Adhatarao, Jiachen Chen, Mayutan Arumaithurai, Xiaoming Fu, and K. K. Ramakrishnan. Comparison of Naming Schema in ICN. In *The 22nd IEEE International Symposium on Local and Metropolitan Area Networks (LANMAN)*, 2016.

[2] Mayutan Arumaithurai, Jiachen Chen, Edo Monticelli, Xiaoming Fu, and Kadangode K Ramakrishnan. Exploiting icn for flexible management of software-defined networks. In *Proceedings of the 1st international conference on Information-centric networking*. ACM, 2014.

[3] Rainer Böhme, Nicolas Christin, Benjamin Edelman, and Tyler Moore. Bitcoin: Economics, technology, and governance. *The Journal of Economic Perspectives*, 29(2):213–238, 2015.

[4] Dan L Burk. Privacy and property in the global datasphere. *Minnesota Legal Studies Research Paper*, (05-17), 2005.

[5] Jiachen Chen, Mayutan Arumaithurai, Xiaoming Fu, and K. K. Ramakrishnan. G-COPSS: A Content Centric Communica- tion Infrastructure for Gaming. In *ICDCS*, 2012.

[6] Jiachen Chen, Mayutan Arumaithurai, Lei Jiao, Xiaoming Fu, and K. K. Ramakrishnan. COPSS: An Efficient Content Oriented Pub/Sub System. In *ANCS*, 2011.

[7] Sergej Dechand, Dominik Schürmann, TU IBR, Karoline Busse, Yasemin Acar, Sascha Fahl, and Matthew Smith. An empirical study of textual key-fingerprint representations. In *USENIX security 2016*.

[8] A. Y. Ding, Bo Han, Yu Xiao, Pan Hui, A. Srinivasan, M. Kojo, and S. Tarkoma. Enabling energy-aware collaborative mobile data offloading for smartphones. In *2013 IEEE International Conference on Sensing, Communications and Networking (SECON)*, pages 487–495, June 2013.

[9] Aaron Yi Ding, Yanhe Liu, Sasu Tarkoma, Hannu Flinck, Henning Schulzrinne, and Jon Crowcroft. Vision: Augmenting wifi offloading with an open-source collaborative platform. In *Proceedings of the 6th International Workshop on Mobile Cloud Computing and Services*, MCS '15, pages 44–48, New York, NY, USA, 2015. ACM.

[10] Yevgeniy Dodis, Jonathan Katz, Leonid Reyzin, and A. Smith. Robust fuzzy extractors and authenticated key agreement from close secrets. *Advances in Cryptology-CRYPTO 2006*, pages 232–250, 2006.

[11] Yevgeniy Dodis, Rafail Ostrovsky, Leonid Reyzin, and Adam Smith. Fuzzy extractors: How to generate strong keys from biometrics and other noisy data. *EUROCRYPT 2004*, pages 79–100, 2004.

[12] Peter Eckersley. How unique is your web browser? In *International Symposium on Privacy Enhancing Technologies Symposium*, pages 1–18. Springer, 2010.

[13] David H Flaherty. On the utility of constitutional rights to privacy and data protection. *Case W. Res. L. Rev.*, 41:831, 1990.

[14] Scott D Halpern, Peter A Ubel, and David A Asch. Harnessing the power of default options to improve health care. *New England Journal of Medicine*, 357(13):1340–1344, 2007.

[15] Tanya Ignatenko and Frans M. J. Willems. Information Leakage in Fuzzy Commitment Schemes. In *IEEE Transactions on Information Forensics and Security*, volume 5, page 337, June 2010.

[16] Ari Juels and Madhu Sudan. A Fuzzy Vault Scheme. *Proceedings of IEEE Internation Symposium on Information Theory*, page 408, 2002.

[17] Ari Juels and Martin Wattenberg. A Fuzzy Commitment Scheme. *Sixth ACM Conference on Computer and Communications Security*, pages 28–36, 1999.

[18] Ari Juels and Martin Wattenberg. A fuzzy commitment scheme. In *Proceedings of the sixth ACM conference on computer and commmunications security*, 1999.

[19] Geoffrey Lightfoot and Tomasz Piotr Wisniewski. Information asymmetry and power in a surveillance society. *Information and Organization*, 24(4):214–235, 2014.

[20] Kirsti Lonka. Promoting flourishing and elevated thought–reflections on e. saarinens pedagogy. *sa aarinen SE*, page 151, 2013.

[21] Tobias Markmann, Thomas C. Schmidt, and Matthias Wählisch. Federated End-to-End Authentication for the Constrained Internet of Things using IBC and ECC. In *Proc. of ACM SIGCOMM, Poster Session*, pages 603–604, New York, August 2015. ACM.

[22] Lynette I Millett, Batya Friedman, and Edward Felten. Cookies and web browser design: toward realizing informed consent online. In *Proceedings of the SIGCHI conference on Human factors in computing systems*, pages 46–52. ACM, 2001.

[23] Lawrence E Mitchell. Innocent shareholder: An essay on compensation and deterrence in securities class-action lawsuits, the. *Wis. L. Rev.*, page 243, 2009.

[24] Richard S Murphy. Property rights in personal information: An economic defense of privacy. *Geo. LJ*, 84:2381, 1995.

[25] M. Raja and S. Sigg. Applicability of rf-based methods for emotion recognition: A survey. In *2016 IEEE International Conference on Pervasive Computing and Communication Workshops (PerCom Workshops)*, pages 1–6, March 2016.

[26] A. Sahai and B.Waters. Fuzzy identity-based encryption. *Advances in Cryptology - EUROCRYPT 2005*, pages 457–473, 2005.

[27] Walter J. Scheirer and Terrance E. Boult. Cracking fuzzy vaults and biometric encryption. In *Proceedings of Biometrics Symposium, Baltimore, USA*, 2007.

[28] Dominik Schuermann and Stephan Sigg. Secure communication based on ambient audio. *IEEE Transactions on mobile computing*, 12(2), 2013.

[29] Paul M Schwartz. Property, privacy, and personal data. *Harvard Law Review*, pages 2056–2128, 2004.

[30] Stephan Sigg, Dominik Schuermann, and Yusheng Ji. Pintext: A framework for secure communication based on context. In *Proceedings of the Eighth Annual International ICST Conference on Mobile and Ubiquitous Systems:Computing, Networking and Services (MobiQuitous 2011)*, 2011.

[31] F Stojano and R Anderson. The resurrecting duckling: Security issues for wireless ad hoc networks. In *Seventh International Workshop Proceedings, Lecture Notes in Computer Science*, 1999.

[32] Pim Tuyls, Boris Skoric, and Tom Kevenaar. *Security with Noisy Data*. Springer-Verlag, 2007.

[33] M. Wagner, J. Laverman, D. Grewe, and S. Schildt. Introducing a harmonized and generic cross-platform interface between a vehicle and the cloud. In *2016 IEEE 17th International Symposium on A World of Wireless, Mobile and Multimedia Networks (WoWMoM)*, pages 1–6, June 2016.