ISSN 2186-7437

NII Shonan Meeting Report

No. 2018-6

RESILIENT MACHINE-TO-MACHINE COMMUNICATION

http://shonan.nii.ac.jp/seminar/114/

Stephan Sigg
Mayutan Arumaithurai
Toru Hasegawa

March 26 - 29, 2018

National Institute of Informatics
2-1-2 Hitotsubashi, Chiyoda-Ku, Tokyo, Japan

RESILIENT MACHINE-TO-MACHINE COMMUNICATION

http://shonan.nii.ac.jp/seminar/114/

Organizers:
Stephan Sigg (Aalto University, Finland)
Mayutan Arumaithurai (University of Goettingen, Germany)
Toru Hasegawa (Osaka University, Japan)

March 26 - 29, 2018

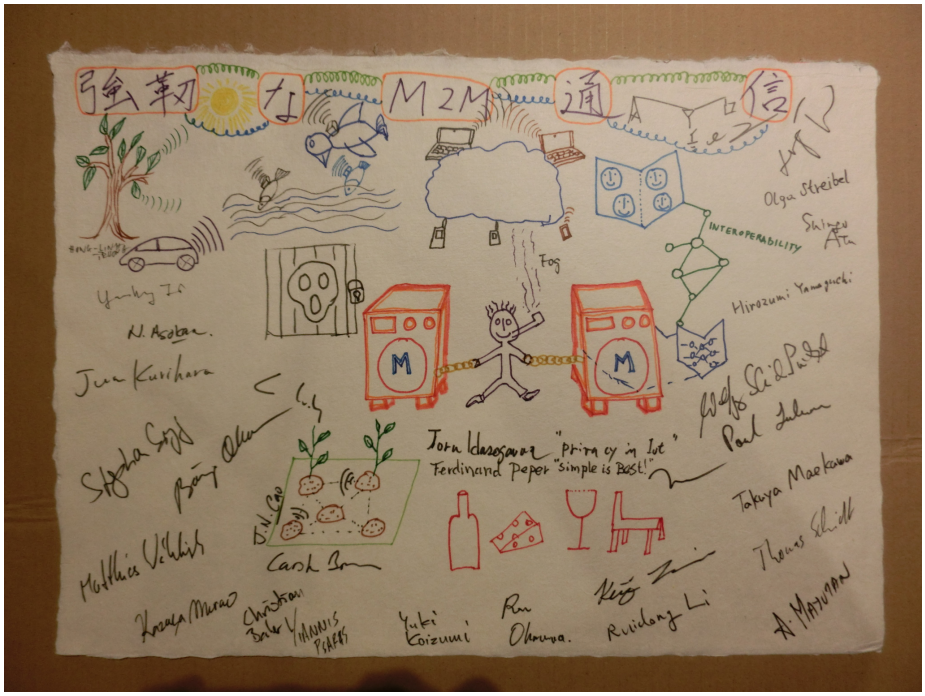

1 Focus of the meeting and Rationale

Machine-to-machine (M2M) interactions such as wearables, vehicular networks and smart homes will constitute more than a third of the total internet connections[1]. These Internet of Things (aka Industrial Internet, aka Industry 4.0) 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 Big Data need to be managed, routed and processed efficiently, for instance, supported by the network edge or fog computing concepts.

These tremendous device and data amounts envisioned will be supported by the upcoming 5G wireless systems which shall support data rates of tens of megabits per second for tens of thousands of connections for massive wireless sensor networks and 1 Gb per second simultaneously. This becomes possible through mm-wave communication, flexible spectrum use, massive MIMO and Femtocells.

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. Networking paradigms are witnessing a shift from location oriented networking to content/information orientation (e.g. ICN, NDN,). The cloud concept transforms the Internet to a network of data centers, featuring computer-to-cloud-to-computer interactions.

The potential benefits of combining the massive environmental perception based on M2M with the control power available in upcoming network paradigms is huge, as is the opportunity of number of research issues opened. Some of the pressing research issues are listed below.

1.1 Networking support for M2M

ICN and SDN have been primarily designed for fixed networks but recent work proposes extensions to wireless networking. 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.

1.2 Wireless support for M2M

5G is envisioned to support M2M scenarios with higher data rate and massive device count. Dynamic spectrum sharing will be required to support optimally adapt to different traffic types and highly variable QoS requirements in M2M scenarios. Further research issues regard cloud radio access networks to move RAN functionality to the cloud for on-demand creation of cloud-based virtual mobile networks exploiting NFV. Vehicular communication is envisioned as one major aspect of M2M.

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

1.3 Security and Privacy for M2M

Security and access control are key concerns for M2M and expand also to secure distributed data structures and privacy preserving data distribution schemes as well as usable security for constrained devices. Attribute-based proxy re-encryption could enable group conditioned access control for M2M content, but also Blockchain holds the potential to serve as a secure distributed database across IoT and Edge devices. Another approach to increase privacy on shared encrypted data is homomorphic encryption for M2M networks Forward and Backward secrecy gains increased importance with the duration of an M2M instrumentation.

1.4 Data support for M2M

While the data produced for M2M is exploding, the nature of the data shifts towards multimedia and video content while at the same time its personal link is intensifying covering health and fitness related data, emotions and data from the private domain. Novel sensing modalities and communication means are exploited with RF based recognition, visible light communication and Intra-body communication. Resource restricted M2M devices require resource sharing mechanisms among devices and the network edge. Several technical solutions have been proposed for this but proper incentives for such collaborative approaches are lacking.

2 Meeting Schedule

The meeting was organized over four productive seminar days with a strong focus on presentations by international experts from academia and industry. In addition, Tutorials, working groups and breakout sessions were organized for in-depth discussion on selected topics. The meeting schedule is depicted in figure 1.

3 Overview of Talks and discussions

Mobile Augmented Reality

Prof. Yu Xiao, Aalto University, Finland

Prof. Yu Xiao introduces their mobile AR systems, e.g., an AR navigation system in a supermarket, an AR app for object recognition, an AR gaming system, and an AR navigation system, and main research challenges for realizing the systems in mobile networks, focusing on difficulties in realizing AR with the current technologies, such as bandwidth consumption or computing. She concludes that mobile AR systems, in general, require high frame rate and high resolution image processing and transportation via networks, and this requires high bandwidth and low latency communications. The network throughput/bandwidth and latency is critical for mobile AR applications and those of 5G systems are not sufficient.

	25.03.2018	26.3.2018	27.3.2018	28.3.2018	29.3.2018
7:30		Breakfast	Breakfast	Breakfast	Breakfast
8:00					
8:30				Challenges Talks (20-30min) Networking, wireless, security and data support for M2M	
9:00		Introduction, 2min intro pitches	Tutorials: (1) Narrow-band IoT (2) Adap-tation & Self Awaren.		Summary of the Seminar, collaboration and continuation
9:30					
10:00					
10:30		Break	Break	Break	Break
11:00		Seminar theme, identifying challenges	Future Opportunities Talks (20-30min)	Breakout Discussions (3rd day topics)	Closing Wrap-up and outlook
11:30					
12:00		Lunch	Lunch	Lunch	Lunch
12:30					
13:00		Group Photo	Future Opportunities Talks (20-30min)	Excursion to Kamakura Great Buddah and Hase temple, Hachimangu shrine, sakura path, shopping/souvenir street	
13:30		Applications Talks (20-30min) Networking, wireless, security and data support for M2M			
14:00			Breakout Discussions (2nd day topics and along 3-4 verticals)		
14:30					
15:00	Check-in				
15:30		Break	Break		
16:00		Xiaoming Fu	Tutorials: (3) IoT comm. standards & Security		
16:30		Breakout Discussions (1st day topics and along 3-4 verticals)			
17:00					
17:30			Dirk Kutscher		
18:00		Dinner	Dinner	Banquet	
18:30					
19:00	Welcome Banquet				

Talks

26.03. 13:30	Speaker	Topic
Networking support	Matthias Waehlisch	Name to MAC address mapping in NDN
Wireless support	Yu Xiao	Mobile Augmented Reality
Security and Privacy	Asokan	Privacy preserving oblivious neural network and adversarial Machine Learning
Data support	Keijo Heljanko	Latency-Constency trade-offs in Distributed Databases
27.03. 11:00		
Networking support	Ferdinand Peper	Communication protocols for highly restricted nodes
Wireless support	Riku Jantti	Ambient and Quantum backscatter
Security and Privacy	Lars Wolf	PotatoNet - Real-world condition WSN (Challenges and lessons learned)
Data support	Paul Lukowikcz	Cyber-Groups - Activity Recognition in Groups
28.03. 9:00		
Networking support	Jiannong Cao	Edge Mesh: enabling scalable connectivity and distributed intelligence for IoT
Wireless support	Yusheng Ji	Vehicular edge/cloud
Security and Privacy	Yuki Koizumi	Privacy issues about ICN and IoT
Data support	Takuya Maekawa	Recent studies with relatively new sensing modalities (CSI, IR camera, Sound probing)

Tutorials, working groups and Talks

Tutorial	Carsten Bormann	Constrained application protocol
Interest group	Mayutan Arumaithurai	Blockchain for M2M applications
Tutorial	Christian Becker, Wolfgang Schroeder Preikschat	Adaption and Self-Awareness in M2M for Predictability
Tutorial	Riku Jantti	Narrowband IoT
Talk	Thomas Schmidt	IoT Networking in RIOT
Talk	Xiaoming Fu	Geosocial Networks
Talk	Dirk Kutscher	Research Directions for Industrial IoT and Edge Computing
Interactive	Olga Streibel	M2M at Bayer

Figure 1: Meeting Schedule

Discussion: Current evaluations of the system were conducted on 4G since 5G is still not in provision, so that further evaluation on 5G might be necessary. Several configurations for LTE systems have been considered and even in the best case, LTE is insufficient to support HD sharing. In simulations, VeinsLTE (SUMO and OMnet++) have been utilized and parameters have been chosen that achieved best performance. The simulator does not support LTE D2D though. In the Hololens demonstration, images were fetched from the FrameReader and converted into jpeg images, which are then sent to the server for processing. Coding technologies will be considered in the future. Regarding wireless support for AR, 5G NR and WiGig capacity should be enough to support image transmission even in dense deployments.

The Need for a Name to MAC Address Mapping in NDN: Towards Quantifying the Resource Gain

Prof. Matthias Wählisch, Freie Universität Berlin, Germany

In this talk, we start from two observations. First, many application scenarios that benefit from ICN involve battery driven nodes connected via shared media. Second, current link layer technologies are completely ICN agnostic, which prevents filtering of ICN packets at the device driver level. Consequently, for any ICN packet, interest as well as data, is processed by the CPU. This sacrifices local system resources and disregards link layer support functions such as wireless retransmission. We argue for a mapping of names to MAC addresses to efficiently handle ICN packets, and start exploring dynamic face-based mapping schemes. We analyze the impact of this link-layer adaptation in real-world experiments and quantitatively compare to different configurations. Our findings on processing, reliability, and energy consumptions on constrained devices indicate significant gains in larger networks.

Further Reading:

- P. Kietzmann, C. Gündogan, T. C. Schmidt, O. Hahm, and M. Wählisch, "The Need for a Name to MAC Address Mapping in NDN: Towards Quantifying the Resource Gain," in *Proc. of 4th ACM Conference on Information-Centric Networking (ICN)*. New York, NY, USA: ACM, September 2017, pp. 36–42.

Discussion: Regarding routing tables and maintenance provided at that layer, there are solutions that also interact with lower levels. A key point though is that the mapping should be independent. The mapping functions needed for interest broadcast use link-layer features, only. It must not be intermixed with higher-layer features such as routing tables. For a compromise between broad- and unicast, multicast MAC addresses and their respective handling is doable as well.

For researchers that want to test their protocols, IoT lab is usable in general, however, it might get challenging for certain applications or test scenarios (e.g., fine-grained energy measurements).

Privacy preserving oblivious neural network and adversial Machine Learning

Prof. N. Asokan, Aalto University, Finland

Abstract: Applications of machine learning (ML) are becoming pervasive in all aspects of human endeavor. Security and privacy applications are no exception. The stunning advances in accuracy and performance of ML-based systems underpins this development. Recent research has brought significant improvements to various applications like detecting malware, steering users away from phishing websites, and helping users by inferring sensible security/privacy configurations based on context.

Any successful system needs to consider and defend against possible adversarial interference. ML-based systems are no exception. Recent work has shown how an adversary could influence the training of ML-models, fool ML-models into incorrect conclusions , or extract information by querying ML-models ("model inversion" and "model stealing"). For example, an adversary can manipulate a traffic sign so that the tampering is not evident to humans, but can fool an autonomous car into incorrectly interpreting it (e.g., left-turn into a right-turn). This is called an "adversarial example". Recent research has shown that it is relatively easy to find adversarial examples against a variety of ML-based systems. Consider another example: cloud-based ML-models are increasingly popular. But these require users to reveal their input data to the cloud server; Input data may be sensitive and can be misused for other purposes; for example, a cloud-based malware detection system results in the server learning about all the applications on a user's device. But this information can be used to completely profile the users by, e.g., inferring their gender, income, hobbies, political/religious affiliations etc.

The potential for adversarial behavior also has an opportunity cost: if multiple organizations can pool their sensitive data, they can build more effective ML-models which can benefit all organizations involved, but will be prevented from doing so, if the sensitive data cannot cross organizational boundaries because of regulation or concerns about potential adversarial behavior. Consider two hospitals in different jurisdictions that can build a better diagnosis model by pooling their patient data, but cannot afford to reveal their data to each other.

Depending on where who the adversary is (e.g., data provider, model trainer, model user) and what its target are (e.g., influencing model training, fooling the model, extracting information about the model or the training data), different security and privacy concerns arise. A number of different techniques can be used to defend against adversarial behavior. Advanced cryptographic techniques like multi-party computation and homomorphic encryption can help design techniques allowing mutually distrusting parties to make joint computations without revealing their private input data to each other. Statistical techniques like differential privacy can enable data sharing by provably ensuring that shared data will not lead to unintended information disclosure. Both of these can provide strong guarantees, but are often expensive. Use of hardware-based security mechanisms, like trusted execution environments, can provide efficient solutions the problem of joint computations on private data but require a leap-of-faith in trusting the integrity of the hardware-based security mech-

anism. A promising architectural alternative to centralized model building is federated learning where multiple data owners can jointly build a model without having to disclose their datasets to the others.

Security and privacy concerns in ML-based systems are multilateral. Therefore solutions are multilateral as well. There will be no one technique that addresses all the privacy and security concerns. Other factors like cost of deployment, usability, and performance also need to be accounted for. As is typical in any large system, a suite of techniques need to be used to provide defense-in-depth.

MiniONN, short for minimalistic oblivious neural networks, is an example of recent efforts to address one such problem. We show how to transform an existing neural network model to an "oblivious neural network" model so that it can be used in a cloud-hosted setting to provide prediction services to clients without compromising the privacy of client inputs or the confidentiality of the server's model. MinoONN is more general and is significantly faster than previous approaches for oblivious neural networks. The work was presented at ACM CCS 2017.

- The full version of the paper is at https://eprint.iacr.org/2017/452
- Slides for the talk are available at https://asokan.org/asokan/research/ML-and-security.pdf

Discussion: Machine learning is ubiquitous, and growing 44 percent annually over the next 5 years. With regard to security, it can be used for (1) Access Control and also for (2) Deception Detection. However, just using machine learning is not sufficient. For example, with face recognition, it is not sufficient to have a high recognition rate, because the adversary could try to mimic the face of authorized persons (mask, makeup). In addition, classification can be thrown off easily by adding carefully selected noise, like for example adding 10 percent of noise to an image to a bus classifies it as an ostrich in some machine learning algorithm, or adding 0.07 percent of noise to a panda makes it a gibbon. This kind of things are problems also in relevant applications, such as with traffic signs in autonomous driving. Another example is that iPhone can be unlocked through ultra-sound. A human would not be fooled, but machine learning is.

Distributed machine learning has a problem if there is a malicious actor. This actor may learn information from other actors without their consent. It is even not clear where to search for the adversary and what is his target. In particular, this is problematic if the adversarial actor resides within the machine learning pipeline, e.g. when the input is compromised. The adversary could, however, also be client that tries to model inference (invert model, infer membership). For example, an adversary could try to steal or approximate the model of Google translate. Furthermore, the prediction service could be malicious. When you have to reveal your set of applications, the service provider could abuse this information. Then, the adversary might reside inside the training pipeline. If a model is optimized for two things at the same time, the person providing the training date may not be aware that some of the data is extracted, even though he/she assumes the data is anonymized. Finally, the data owner could be malicious. For example, SPAM mails hold a lot of hidden information that confuses the training of the SPAM filter. Other example is Microsoft's on-line

learning chat-bot that was provided with incorrect/false data, making the model eventually behave in unintentional ways. In summary, if there is a reason for someone to subvert your model, you have a problem. Therefore, cloud-based machine learning might be problematic. A natural question is therefore whether is is possible to make a Neural Network oblivious, i.e., that it is basically unaware of the data. The speaker did this by using an interactive protocol: MiniONN. This proof-of-concept implementation achieved an overhead of 1 second, which may be still too much for real Machine Learning applications that might expect microsecond delays.

Another possibility is, to use machine learning to defend against such attacks, i.e. to distinguish between adversarial and benign behavior. The caveat is that you only get probabilistic guarantees, so that such solutions are potentially susceptible to evasion themselves. Overall, system level defenses are probably the best way, at least in the short-term, to defend against adversarial behavior The holy grail is to understand how humans do inferences and have Machine Learning do it like that. But in the short- to medium-term, we still need to defend current machine learning techniques from adversarial behavior. There are approaches to add noise to data (differential privacy) or warn users when they have revealed enough information to a service to allow leakage of unintended information. Problem with the first is that adding excessive noise may harm utility (i.e., make the resulting dataset not useful for the purpose it was intended for) and with the second is that the adversary may be smarter and could make meaningful inferences with less data than the defender expected. Also adding noise might cost you, e.g. in terms of bandwidth.

Latency-Consistency trade-offs in Distributed Databases

Assoc. Prof. Keijo Heljanko, Aalto University, Finland

Abstract: One of the main approaches to coordinate distributed systems is to use a fault tolerant distributed database to store the data needed for coordination between nodes. The CAP theorem by Brewer has shown that under network partitions such distributed databases can only have two of the following three properties: Consistency (C), Availability (A), and Partition tolerance (P). Thus the user of databases needs to make a concious choice on whether to have a database that is centralized (CA) leading to limited scalability, consistent and partition tolerant (CP) leading to a database that will be consistent but will not be available for writes during network partition, or available and partition tolerant (AP) leading to inconsistencies during network partition. Practical applications are often combinations of CP database systems for data that needs consistency and AP database systems for data that needs low latency. The talk also discusses the use of immutable data and conflict free replicated datatypes (CRDTs) as application patterns allowing the use of AP datastores without problems with data inconsistencies. The use of highly syncronized atomic clocks to do distributed transactions between database shards was discussed as an implementation technique used by the Google Spanner database. This would require reliable high precision clock syncronization, a feature that would be most welcome to also commercial cloud database servers.

Discussion: An interesting detail described in the Google Spanner paper is that it suffices that Google has time synchronization guarantee as high as 6 ms. Unfortunately, commercial cloud vendors are not yet giving high-precision clocks as a service. Such more precise real-time information for applications would allow the same techniques as Google Spanner has implemented using their proprietary TrueTime service to be also used in commercial cloud services.

Further reading:

- Corbett et al. Spanner: Google's Globally Distributed Database. ACM Trans. Compt. Syst. 31(3):8:1-8:22 (2013)

DeepScan: Exploiting deep learning for malicious account detection in location-based social networks

Xiaoming Fu, University of Göttingen

An adversary may give harmful effects for PoI-centric location-based social network applications. For example, malicious users may check-in within minutes to multiple different locations. The presenter focuses on the detection of malicious users in LBSN. Machine learning is used and detecting malicious accounts in Dianping (Chinese version of Yelp) is presented as an example. For data annotations, 15 volunteers are recruited, and more than 1/3 were malicious. The method is based on the analysis of time series activities (e.g. user profiles, check-in information, reviews). DeepScan uses Long-Short-Term-Memory (LSTM) to judge whether an account is malicious or not. Results show high accuracy of 98%. If time-series features are not considered, it becomes 95%. Compared with the other approaches, the proposed one overwhelms them. In summary, by fully-utilizing tempo-spatial features, high accuracy could be achieved.

Discussion: The proposal did not yet consider individual account identification, so that an attacker could use multiple accounts (Sybill attack). The presenter stated that, currently, students are hired to annotate the malicious accounts.

Tutorial: Adaptation and Self-Awareness

Christian Becker, University of Mannheim, Wolfgang Schroeder-Preikschat, University of Erlangen-Nuremberg

Ant colonies are an example of a self-organizing system. In general, self-organizing (Software) Systems feature the five aspects (system state, evaluation criteria, acceptance space, disturbances, control mechanisms). For such systems, system objectives vs. low-level control need to be considered as well as robustness (dead space, survival space, acceptance space, target space), flexibility and limitations of adaptivity (adaptation vs. awareness). In particular, recovery is possible only from a survival space but disturbance might lead into a dead space. A use case for such self-organizing system is platooning, where

different approaches can be used for control: (de)centralization and distribution. Overall, systems become more and more context-aware, communicate, and autonomous. Adaptation and self-awareness are key concepts.

For M2M, always some kind of OS is involved, especially in interprocess communication, and device programming. Several design decisions impact system performance: One stack per instance gives lower latency while one stack per kernel gives higher latency. For sharing common resources among systems process scheduling is one aspect for which we have to come up with a sequence of actions.

Energy efficiency is a major problem: A single bitcoin generation uses lots of energy and therefore large-scale IoT means large-scale energy problems. Predicting energy demand is good but the hardware converts energy however the software determines how much. What can be done is static program analysis (worst case execution analysis) to accumulate knowledge. Essentially, this way we are counting the Joules of a single instruction. An application where this was used is the flying sensor net, a project with biologists and flying bats. They designed a sensor(2 g weight) for the 20 g bats, which consists of a microcontroller with own energy profile. Energy awareness is not only a technical but also an economical and ecological issue.

Discussion: It appears that in computer science, known ideas from control theory are re-invented. What seems missing so far are asynchronous interactions, deadlocks and so forth. The impact of latency is typically severe and network and latency is critical.

Performance of Narrowband IoT

Prof. Riku Jäntti, Aalto University

LoRa and SigFox are examples of technologies for low bandwidth long range communication technologies. They are now being complemented with Narrowband IoT (NB-IoT) which is a non-backwards compatible extension of LTE targeted for cellular based IoT applications. Strong beamforming is used to concentrate energy in the radio to specific target devices. The system is reliable through multiple retransmissions.

Two major classes of IoT communication are Massive IoT & Critical IoT. The main difference is that Critical IoT is time critical. The goal is to allow for up to 10 years battery lifetime for NB-IoT devices, depending on usage pattern. For comparison: SigFox is uplink only without any downlink, no acknowledgements. LoRa has some downlink but is mainly focused on uplink.

Aalto has implemented a NB-IoT prototype. They have a base station and IoT module implemented in a software defined radio. All the base band processing in the base station is done in a Linux PC. Regarding virtualization, NB-IoT is easier to virtualize than regular LTE as there is less front-haul traffic. The low bandwidth also reduces the time needed for computation, which allows longer distances to computing resources or slower computers.

It seems clear that NB-IoT will actually be deployed. Basically all operators have deployment plans and it is already deployed in China.

Discussion: Chip size of NB-IoT devices can be 26x16x2.5 mm and PHY layer performance is collision-free (different to LoRA). It also scales much better than LoRA (https://tools.ietf.org/html/draft-ietf-lpwan-overview).

Further reading:

- M. R. Palattella et al., "Internet of Things in the 5G Era: Enablers, Architecture, and Business Models," in *IEEE Journal on Selected Areas in Communications*, vol. 34, no. 3, pp. 510-527, March 2016.
- Comparison of Wi-Sun with other technology, done by Wi-Sun alliance: https://www.wi-sun.org/index.php/tcwp-en/file

Communication protocols for highly restricted nodes

Ferdinand Peper, NICT-CiNet

Possible applications of neural dust are medical and intelligent materials. Medical applications are wearable, implantable devices. For instance, a dust device in brain can communicate to the external device attached to the skull. Neural dust is also used for muscle control. Regarding intelligent materials, biological systems tightly integrate sensing, actuation, and control. Also, engineering applications exits that could benefit from a similar approach.

Operating conditions for high-density sensor networks are *Energy-autonomity* (make nodes simple, make operation simple, use energy-efficient signaling), *Expendability* (low cost, flexible), *Ubiquitous* (high density. small size, low cost), *Non-interference* (wireless range small, difference modes of wireless). Therefore, in the group of the presenter, spiking is used for low-power consumption. The general philosophy is that there are no node IDs, no routing tables, no multi-hop routing, not device-oriented, but location-oriented. For instance: communication through silence (Interval of spike is used).

Discussion: The devices are operating in MHz frequency and the propagation delay of two spikes are identical. The nodes, however, never know that they reached consensus.

Ambient and Quantum Re/Back-scatter Communications

Prof. Riku Jäntti, Aalto University

Backscatter of radio waves from an object has been a subject of active study since the development of radar back in the 1930s, and the use of backscattered radio for communications since Harry Stockman's work in 1948. Backscatter Communications (BC) is widely used in RFID where a reader device generates an unmodulated carrier signal, a passive tag absorbs the energy of this signal and then sends back the modulated signal to the reader. BC devices do not need a power-hungry transceiver and can achieve up to 1000 times lower power consumption and 10 to 100 times lower device cost than contemporary active-transceiver-based solutions. In traditional BC solutions, a reader device needs to spend power transmitting unmodulated carrier that will be then modulated by

the BC device. In Ambient backscatter communications, the BC devices modulates the ambient signal impinging at their antenna. It reuses the power and radio spectrum of other wireless systems to transmit it's information without causing harmful interference to these systems.

Microwave quantum technology is becoming more mature as more and more components have been demonstrated in laboratory environments. One of the foreseen applications of microwave quantum technology is the quantum radar (QR). Similarly as backscatter communications bear close resemblance to classical radar technology, the Quantum Backscatter Communications is closely related to the QR. In QBC, quantum phenomenon are utilized to improve the system performance beyond the physical limits of their classical counterparts.

Discussion: In backscatter communications you pay the price in the ambient sender power consumption. You have to use some coding to detect backscatter signals from multiple devices. What would help though is to know the sent FM signal (e.g. receiving the network radio broadcast). In theory, quantum backscatter can be implemented but it needs cooling to a very low temperature (mK level). It would be needed to reach ultimate sensitivity by minimizing the number of photons detected.

PotatoNET - Real-world condition WSN

Lars Wolf, Technische Universiät Braunschwig

Our PotatoNet has been deployed on an agricultural area in 2015 to perform several WSN outdoor experiments while measuring the stress of potato crops. It was extended a year later by the PotatoMesh, a solar panel-based mesh network of nodes, and another deployment in 2017. Throughout all these deployments we experienced problems and failures at different stages of the projects. We derive key problems and some important concepts when it comes to outdoor WSN deployments. Also the storing of the crop is an important part of the overall agricultural production and logistics chain. Monitoring and adjusting the environmental conditions in storage are important tasks to achieve an optimal and efficient storability. Therefore, we are currently investigating the communication part of such a system where information about current conditions of stored potatoes can be collected from sensor nodes which are in the middle of stored potatoes. Among others, we study suitable frequencies and networking structures.

Discussion: Agriculture is an important application domain for WSNs. The battery should optimally last for at least one season, the measurement frequency was roughly 1 sample per minute. The voltage scaling was made fail save by using a separate module and by internal checking (matrix computations).

Further reading:

- Jakob Juu, Wireless Sensor Networks and Localization for Biomass Storages, Aarhus University, Denmark 2015.

Cyber Borg and Cyber-Groups: Activity recognition in groups

Paul Lukowicz, German Research Center for AI, DFKI

Access to digital domain is increasingly more frequent and becomes the center of our life: any time any place. Positive aspects that evolve from this are e.g. Assistants that are listening all the time, dietary monitoring by classifying chewing sounds, smart phones predicting divorces. We are at the verge of any real life event becoming, instantly globally connected to the digital domain. Anything that happens in physical life leaves a digital imprint. Physical objects have digital shadows. You can use devices to personalize your spam and to interact with the digital shadow

Discussion: A great application would be noise-cancelling headphones to cancel advertisements – one of Steve Mann's favourite applications: alternative reality. Regarding privacy, the right balance is important. Neither 'forget about privacy' nor putting people with flags in front of cars is appropriate. In any case, companies already know details about people. Health insurance company can know your health data because you can trust them – they would go bankrupt if they handled your data like Facebook. We do have social norms that support here too. Comment: Privacy violations by careless third parties. GDPR – 'your data can be moved outside of the EU'.

Research Directions for IoT Edge Computing

Dirk Kutscher, CTO Huawei Munich (Presented remotely)

Computing is shifting to distributed perspectives. Industrial IoT use cases include TSN/Profinet/CAN which require video feeds, real-time analytics and real-time control loops processing. IoT data streams are directed to the cloud from many IoT domains (Smart City IoT, Industrial IoT, Home IoT). Current in-network computing with C/S protocols has many limitations (scalability, efficiency, performance and robustness). To address the challenges, Unikernels, Light-weight scripting, Trusted execution environments, Data-oriented communications and programming abstracts (ICN/NFN), Distributed consensus protocol (blockchains), are studied.

NFV adds computation to networks. Networked computation is a consequence of adding more and more to the network leading networks to become distributed systems themselves.

Data silos can be established and should be avoided. This can be hardened by the constant flow of data from sensors and IoT applications. Named Data Networks are described as a client server infrastructure where the name spaces are managed by servers. Obviously C/S architectures have some problems here, eg. firewalls, mapping to underlying L2 protocols. As a consequence, networks should be built where computation is a first-order service.

Opportunities for ICN were sketched and the link to IoT was highlighted. An extension of ICN is named function networking. Similar to stored procedures NFN can be used for data intensive applications.

Named functions can be used to virtualize execution. In contrast to a RPC the network decides the placement. This is complemented by SPOC, a protocol to play for decentralized computations as nano payments.

Opportunities for NFN research are, e.g., transport based on MQTT, client authentication, auto-scaling, error semantics and QoS, and abstractions for programming are needed.

Summarizing, challenging topics in ICN/NFN research are

ICN security, opportunistic caching and in-network caching, pub/sub, custodian storage/forwording/processing, multi-tenancy.

Named function networking dynamic computation (for edge computing big data, streaming processing, service chaining, Broker function, scalable computation (auto-scaling), efficiency, deterministic computation (performance aspects), programming models).

Some related pointers to further reading:

- Sifalakis et al.: An Information Centric Network for Computing the Distribution of Computations, 1st International ACM Conference in Information Centric Networking (ACM ICN 2014), September 2014, Paris, France.
- Krol et al.: NFaaS: Named Function as a Service. In Proceedings of the 4th ACM Conference on Information-Centric Networking (ICN '17). ACM, New York, NY, USA, 134-144.
- Krol et al.: "SPOC: Secure Payments for Outsourced Computations", NDSS'18 Workshop on Decentralized IoT Security and Standards (DISS).
- https://trac.ietf.org/trac/irtf/wiki/icnrg
- https://datatracker.ietf.org/rg/dinrg/about/

Discussion: Towards the future of edge computing, the environment is dynamically changing. Potential players are going into good positions, e.g., the Mobile China approach. Here, chances for mid sized enterprises exist. Edge computing can be used for mission critical applications, such as vehicular networks, but problems exist with extensibility of car functionality. Here, good separation of computing can help.

Machine learning

Prof. Paul Lukowicz, German Research Center for Artificial Intelligence, DFKI

The basic Machine learning approach is to find features, collect data and find separation boundary between classes. The underlying assumption is that there is a correspondence between the structure of data and semantics. There is also an assumption of smoothness: if a test data point is closer to a training data point, it will be classified correctly. The goal of machine learning then is to find a computational form of the mapping from feature space to class space.

The naive approach is to try all possible functions. Machine learning retains uncertainty (there can be several possible boundaries between classes; this is

what leads to the existence of adversarial examples) and will have to model noise. This is essentially a search problem. With infinite time/memory, and noise-free training data it is possible to build a perfect model. The goal of practical machine learning is to find a reasonable approximation with finite data and finite resources. The underlying approach to train a model is to find a family of parametrized functions, define an error function, and find values for parameters that minimize the error. Different approaches to do this lead to different training algorithms: decision trees, linear classifiers (separated by a plane), perceptrons (weighted inputs, sum, and apply to an activation function). Fundamentally this is a multi-dimensional minimization problem that is not solvable analytically. Machine learning tries to find approximations. The number of parameters should be minimized because more parameters implies the need for more training data. A simple approach is to use one parameter set per region (every region has its own parameter). Shared parameter sets, where the same parameter can differentiate between multiple classes are more powerful. This enables the use of exponentially fewer parameters. Deep neural networks work by hierarchical representation: e.g., in handwriting recognition, intuitively, the first layer discovers shapes in fixed positions, the next layer does position-independent shape detection (self-organized feature detector layers). Nested representation can capture any manifold with distributed representation. The renaissance of deep neural networks is due to two reasons: (1) people didn't know how to train deep networks, such that local minima can be avoided and (2) the computational complexity. The breakthrough came with the realization that training layer by layer first is the key (unsupervised training of autoencoders that build hidden layers as compact representations allowing restoration of data without loss; similar to dimensionality reduction but without constraints).

Discussion: Message/Take away: Machine learning algorithms are by no means intelligent or even really learning. It is more that a function is optimized with respect to some underlying data and there is no real intelligence or learning in this. In popular culture and also by popular people, the 'risk' of machines taking over due to Machine learning is drastically exaggerated. It is not comparable with the abstraction capabilities that humans have: Show a human a single picture of a cat and he can abstract from this and knows how cats look like. This is by no means possible with today's Machine learning approaches.

On the other hand, in popular text mining literature, it is pointed out that the approach taken by Machine learning can indeed be described as learning: We show examples and train the system on these examples and from these examples, the recognition function becomes better optimized. This is, on some level of abstraction similar to the human learning process.

Implementing a feedback loop for Machine Learning approaches might in the optimum case lead to the construction of a turing machine. This is also related to control theory and feedback control loops in that discipline. Indeed, the Machine learning approaches are running on a Turing machine, so that is the most capability we can expect from them. It is impossible to achieve any computation model superior to a Turing machine when this was the computation model where the algorithm was executed on.

Edge Mesh: Enabling scalable connectivity and distributed intelligence for IoT

Prof. Jiannong Cao

Smart IoT and edge computing can help to make IoT smarter. Edge computing research can be divided into parts: Interaction with IoT devices; Architecture and Management. Most research is focused on single edge nodes. Edge mesh is, a solution to leverage mesh network and system architecture of edge devices to enable collaboration and integration of edge devices to support large scale IoT.

Edge mesh scenarios constitute connecting machines in a construction site, machines in an industry, vehicular networks, healthcare, intelligent eco system. The design principles applied cover decentralized architecture, bottom up coordination, multi hop connectivity, gateway for interoperability, pushing intelligence from cloud to edge devices. This can create a distributed intelligence where Ede devices collaborate for distributed decision making. The presenter has distributed intelligent SDN controllers (Distributed task computation; How to assign tasks to distributed edge nodes). Future work constitutes extending SDN for M2M communications, device mobility, Edge mesh for 5G.

Discussion: To set up a testbed for Edge mesh, the presenter started from a wireless mesh networks and used and added edge devices to connect heterogeneous networks. They integrated the edge devices into these wireless networks. Integrating vehicular networks creates security and privacy issues. Another question is whether the cloud is actually needed with so many edges in the future. There might be a situation where the application switches between centralized and distributed.

Vehicular Edge/Cloud: Offloading Framework & Scheduling

Prof. Yusheng Ji, National Institute of Informatics, Japan

Considering the future of smart vehicles and their high computing demand, this talk treats the problem of job assignment and offloading to neighboring vehicles. To solve this problem, the speaker introduced two computing schemes: AVE (Autonomous vehicular edge) and HVC (Hybrid vehicle cloud) schemes. For AVE, no external infrastructure (like cloud servers) is used and jobs generated at a vehicle are directly offloaded to its neighbor vehicles with job scheduling based on Ant-Colony optimization (ACO). The speaker also showed that the ACO rapidly converges. For HVC, external infrastructure, i.e., road side units (edge nodes) and cloud servers are implemented to support the job offloading to neighbor vehicles and to meet the real-time requirements to computation. Due to the uncertainty of future job arrival and resource availability, online scheduling is adopted and job queueing in AVE is removed. The speaker also described that thanks to the location estimation with GPS, the success rate of job offloading (successful finish of offloaded jobs) increases.

Discussion: The computation is performed in an on-line manner or by using the ACO-based algorithm and the time complexity of the scheduling is polynomial in the number of jobs, and the number of nodes. It hence completes the computation in a short time. The problems considered by the speaker include that beacons are not relayed, and offloading requests can only be forwarded till the second hop. So the problem is not very large, i.e., two hops at most.

Privacy Issues in ICN and IoT

Prof. Yuki Koizumi, Osaka University, Japan

Information Centric Networking (ICN) inherently provides many benefits that IP cannot offer, such as in-network caching, client mobility, multicast, and security, thanks to its name-based forwarding, i.e., packets are forwarded/routed according to application-meaningful names. Despite the benefits of ICN, privacy in ICN may be harder than that in IP because ICN uses meaningful names of packets, which are visible to all forwarders. An existing study has proposed *name obfuscation*, which encrypts meaningful application data names into obfuscated meaningless network names and sends Interest packets to the obfuscated names. Although name obfuscation may mitigate the threat, a certain degree of information might be leaked if attackers use auxiliary information, such as popularity of content. The speaker has proposed a method based on k-anonymity and l-diversity to mitigate the threat to privacy in ICN. With the proposed method, a consumer sends an Interest packet with $k-1$ dummy Interest packets, which have l diverse properties, so that attackers can neither distinguish the Interest packet of the consumer from the other $k-1$ Interest packets nor estimate a certain degree of information from the Interest packets.

Discussion: The Key point of ICN is that it breaks up the information monopoly of Google and the likes. Names of data pieces are stored together with the data. Data pieces are encrypted so that only authorized consumers can decrypt them by using any of existing cryptography techniques, such as public-key cryptography, and hence they cannot be seen by attackers. Second, it is difficult to get names by attackers because they are encrypted. That is, it is difficult to get data pieces in the cache by specifying the names of the data pieces.

Recent Studies with Relatively New Sensing Modalities

Prof. Takuya Maekawa, Osaka University, Japan

Several activity recognition projects have been presented. One example is state change detection by using Wi-Fi CSI. The goal of the project is to monitor elderly people at home in a non-invasive way. Only a single receiver is used, and movement of objects in a room, such as doors, furniture and windows, is detected. CNN is used, and a high f-measure was achieved (>0.8). Also several indoor/outdoor positioning methods are introduced. From them, a method to estimate indoor location semantics. The idea is to identify the location features by smartphone sensors and active sound probing (wall material classification). The accuracy was not sufficiently high. Finally, a bio-navigation project is

presented. The issue is to save energy of video camera attached to animals, and the idea is to trigger video recording (and GPS) by accelerometers. In field experiment, sensor loggers are attached to wild birds in Japan, and the video from bird-eye view is presented. They are also attached to cormorant.

Discussion The speaker trained the model using all the data and test the model by cross validation.

A dozen years of standardizing the Internet of Things

Prof. Carsten Bormann, University of Bremen, Germany

The slides of the talk are available here: https://drive.google.com/open?id=14umUCrsYZHbZHCZDwik7ibrXY7rgKicE

Networking the IoT with RIOT

Thomas C. Schmidt, HAW Hamburg, Germany

The Internet of Things (IoT) is rapidly evolving from large numbers of embedded devices that gradually connect to the Internet. Such nodes are often constrained and limited to battery-powered low power lossy radio links. RIOT, the friendly operating system for the IoT, is an open source initiative for fueling an IoT ecosystem that is not locked in with vendors or service operators.

This talk introduces the networking architecture that turns RIOT into a powerful IoT system, and enables low-power wireless deployment. RIOT networking offers (i) a modular architecture with generic interfaces for plugging in drivers, protocols, or entire stacks, (ii) support for multiple heterogeneous interfaces and stacks that can concurrently operate, and (iii) GNRC, its cleanly layered, recursively composed default network stack. Focussing on deployability, we discuss and analyse several IoT networking approaches including 6LowPan and Information Centric Networking.

Further Reading:

- M. Lenders, P. Kietzmann, O. Hahm, H. Petersen, C. Gündoğan, E. Baccelli, K. Schleiser, T. C. Schmidt, and M. Wählisch, "Connecting the World of Embedded Mobiles: The RIOT Approach to Ubiquitous Networking for the Internet of Things," Open Archive: arXiv.org, Technical Report arXiv:1801.02833, January 2018.

4 Overview of Interest Group Discussions

In break-out sessions, we have worked on selected topics that evolved from the discussions. The results are briefly summarized in the following.

4.1 Break-out topic 'Machine Learning'

4.1.1 Group 1

- Machine Learning is not about learning, it is about multi-dimensional data fitting
- data fitting benefits from additional data points and plausability checks; quality of input is essential
- Humans would like to understand why ML-based decisions were taken (and in which manner). For control theory approaches, stability was always important, here is significant difference (at the moment?).
- we now have a lot of systems that sense data and use machine learning; they may contribute to a global view (and go from lower layer data to higher layer semantics)
- assume that every object can communicate with every object in a cheap way, we can easily enable such a global view
- you can combine these sources, this combination (fusion) can provide for a kind of sanity check (of ML decisions)
- this interconnection is a game changer, similar to atomic clock in distributed systems domain
- low latency communication is one requirement to enable specific application scenarios
- trust problems arise when you incorporate additional 'sensors'
- global view could lead to information overload (edge computing may help?)
- selection of information to be provided to others probably needed, also to avoid network congestion

4.1.2 Group 2

- Potential example applications of machine learning are traffic control, attack detection, autonomous vehicles, production/maintenance of smart grid systems and so on.
- Among them, autonomous vehicles have attracted attention. Resiliency against malicious vehicles is important.
- Identifying communication patterns of IoT devices and attackers is hopeful.
- Distributed machine learning is a good example of edge computing. Protecting privacy is its advantage.
- Edge computing (machine learning) under unreliable wireless networks is difficult.
- The other issues include bad effects due to compromised machine learning nodes, roles of humans, feasibility of brain computing, learning based on simulations and so on.

4.1.3 Group 3

- We discussed if machine learning can provide guarantees, similar to real time systems. We discussed how airplanes can provide good guarantees, but at a high cost and high redundancy. Eventually, what we arrived at is that what we can provide is confidence pertaining to certain environmental

conditions. E.g., an autonomous train could provide higher confidence soon since the environment is better controlled.

- Metrics: We are probably looking at the wrong metrics or comparison. E.g., probably machines cannot detect an apple with 100% guarantee, but can probably detect an apple while you are drawing due to the amount of data that they have access to. For instance, take a look at quickdraw (https://quickdraw.withgoogle.com/) and its database (https://quickdraw.withgoogle.com/data).
- when can machines start to program? That might take a while, however, it might be interesting to see if machines can automatically generate documentation by looking at the code written.

4.2 Break-out topic 'IoT'

We first structured the topic into (Radios (PHY/MAC); Networking (IP, ICN); Application Layer/Transfer; Structural interop/serialization; Semantic interoperability; Data aggregation, machine learning). In parallel to these building blocks, we have security and privacy aspects. Naming, discovery, and self-description relate sometimes to protocols, sometimes not. With discovery we mean neighbor discovery (lower layer) but also service discovery etc. One open question is where computation is performed, but we did not discuss this further. In the second part of this breakout, we focussed on different directions in radio development. We distinguished between (i) wide area radio and (ii) ultra low resource radio.

Wide area radio We discussed in more detail LoRA, SIGFOX, WY-SUN, NarrowBand IoT. SIGFOX is considered a very proprietary service, which is the reason why many people started to consider LoRa. We agreed that you cannot build a full network based on SIGFOX, but you may want to support some kind of client functionality. The range of radio depends on the deployment details. In general, there are two ways of improving coverage, either by a very good radio (power strength), or by designing a cellular model. Then, we discussed the problem of coexistence of different administrative radio domains (see work by Laura Feeney). These different domains might be because of different radio technologies but might also exist when the same technology is deployed. Severeness of the problem depends on upcoming IoT deployment. We agreed that further research should be conducted to avoid future problems. We also discussed support of IP layer convergence for the different radio technologies. LoRa, NB-IoT, and WY-SUN already provide IP layer convergence, SIGFOX is getting there slowly (main probleme is very small packet sizes). Basically, for all wide area technologies, an IP convergence layer can be designed. This is in contrast to ultra low resource radios, for which we do not know how to design a convergence layer.

Ultra low resource radio We identified LoRa on backscatter, Inter-only interval, BRZZ (work by Lars Wolf), power line communication, and ultrasound communication. Most of these technologies are more on the research side. After discussing radio technologies, we briefly discussed the networking layer, including IP, ICN, DTN, and deterministic networking. We highlighted

that IP can be used in very different ways. In particular in the IoT, where we have very heterogenous link layers, tweaking IP can be useful to improve they way how networks are constructed (e.g., bootstrapping, management). Further details on this topic will be presented in the tutorial given by Carsten Bormann and in the talk given by Thomas Schmidt.

4.3 Break-out topic 'Blockchain for M2M'

We discussed whether blockchain is really useful for M2M/IoT networks, where we can use it in M2M/IoT, and why we use it from perspectives of its heavy and slow computation and requirements on high computing resources. The following issues/points arises during the discussion:

- To circumvent the heavy computation of blockchain, constructing a hierarchical structure, where blockchain is not applied to check transactions among IoT devices in the lower layer and it is applied to check groups of the transactions in higher layers.
- The blockchain technology can prove that registered transactions are surely carried but it cannot prove that data pieces provided to the transactions are surely generated by authorized IoT devices.
- Blockchain is useless if all data of IoT/M2M systems is gathered in the cloud. That is, blockchain can be used to disconnect the relation of the cloud and the IoT/M2M systems.
- Inserting a block into blockchain might be fast enough for registering new trusted IoT devices, and hence blockchain can be applied for managing IoT devices instead of managing each transaction between the devices.

Finally, we have concluded that blockchain can be used to prove trusts of nodes participating in M2M/IoT systems, to retrieve reputations of the nodes, and to realize incentives/rewards for nodes to supporting the distributed system.

5 List of Participants

1. Ioannis Psaras, UCL, UK, (i.psaras@ucl.ac.uk)
2. Thomas C. Schmidt, HAW Hamburg, Germany (thomas.schmidt@ids-mannheim.de)
3. Matthias Wählisch, Freie Universität Berlin, Germany (waehlisch@ieee.org)
4. Lars Wolf, Technische Universität Braunschweig, Germany, (wolf@ibr.cs.tu-bs.de)
5. Carsten Bormann, Universitt Bremen, Germany (cabo@tzi.org)
6. Hong-Linh Truong, TU Wien, Austria, (hong-linh.truong@tuwien.ac.at)
7. Ruidong Li, National Institute of Information and Communications Technology (NICT), Japan (lrd@nict.go.jp)
8. Jun Kurihara, Zettant Inc., Japan, (kurihara@ieee.org)
9. Yuki Koizumi, Osaka University, Japan, (ykoizumi@ist.osaka-u.ac.jp)
10. Prof. Shingo Ata, Osaka City University, Japan (ata@info.eng.osaka-cu.ac.jp)
11. Prof. Christian Becker, University of Mannheim, Germany (christian.becker@uni-mannheim.de)

Figure 2: Participants of the Shonan Seminar 114 (Resilient Machine-to-Machine Communication)

12. Prof. Jiannong Cao, The Hong Kong Polytechnic University, Hong Kong (jiannong.cao@polyu.edu.hk)
13. Prof. Keijo Heljanko, Aalto University, Finland (keijo.heljanko@aalto.fi)
14. Prof. Riku Jntti, Aalto University, Finland (riku.jantti@aalto.fi)
15. Yusheng Ji, National Institute of Informatics, Japan, (kei@nii.ac.jp)
16. Prof. Paul Lukowicz, German Research Center for Artificial Intelligence, DFKI , Germany (paul.lukowicz@dfki.de)
17. Prof. Kazuya Murao, Ritsumeikan University, Japan (murao@cs.ritsumei.ac.jp)
18. Prof. Jrg Nolte, Brandenburg University of Technology, Germany (jon@informatik.tu-cottbus.de)
19. Dr. Borje Ohlman, Ericsson Research, Sweden (borje.ohlman@ericsson.com)
20. Prof. Ren Ohmura, Toyohashi University of Technology, Japan (ren@tut.jp)
21. Dr. Olga Streibel, Bayer, Germany (olga.streibel@bayer.com)
22. Prof. Xiaoyan Wang, Ibaraki University, Japan (wangxy@nii.ac.jp)
23. Prof. Yu Xiao, Aalto University, Finland (yu.xiao@aalto.fi)
24. Prof. Hirozumi Yamaguchi, Osaka University, Japan (h-yamagu@ist.osaka-u.ac.jp)
25. Prof. Xiaoming Fu, University of Goettingen, Germany (fu@cs.uni-goettingen.de)
26. Prof. Takuya Maekawa, Osaka University, Japan (maekawa@ist.osaka-u.ac.jp)
27. Prof. Ferdinand Peper, NICT-CiNet, Japan (peper@nict.go.jp)
28. Prof. Wolfgang Schrder-Preikschat, Friedrich-Alexander-Universitt Erlangen-Nrnberg, Germany (wolfgang.preikschat@t-online.de)