Report from Dagstuhl Seminar 23222

Novel Scenarios for the Wireless Internet of Things

Haitham Hassanieh*[1], Kyle Jamieson*[2], Luca Mottola*[3], Longfei Shangguan*[4], Xia Zhou*[5], and Marco Zimmerling*[6]

1 **EPFL – Lausanne, CH.** haitham.alhassanieh@epfl.ch
2 **Princeton University, US.** kylej@princeton.edu
3 **Polytechnic University of Milan, IT.** luca.mottola@polimi.it
4 **University of Pittsburgh, US.** longfei@pitt.edu
5 **Columbia University – New York, US.** xia@cs.columbia.edu
6 **TU Darmstadt, DE.** marco.zimmerling@tu-darmstadt.de

Abstract

The Internet of Things (IoT) aims to network everything near and far in our ambient environment. Although the functional innovations for IoT are going full steam ahead, newly-emerging scenarios such as the Internet of Ocean and Implantable Things often come with limited power budgets, challenging deployment scenarios, and demanding computational resources, which fundamentally stress conventional IoT architecture, communications primitives, and sensing capabilities. The goal of this Dagstuhl Seminar was to bring together researchers from both academia and industry globally to *i*) review the capacity of existing IoT research from radical perspectives; *ii*) summarize fundamental challenges in modern IoT application scenarios that may then be investigated in joint research projects; and *iii*) discuss new types of hardware architecture, network stack, and communication primitives for these emerging IoT scenarios.

Seminar May 29 – June 1, 2023 – https://www.dagstuhl.de/23222

2012 ACM Subject Classification Hardware → Emerging technologies; Networks → Network architectures; Networks → Public Internet

Keywords and phrases Internet of Ocean Things, Internet of Medical Things, NextG Communication

Digital Object Identifier 10.4230/DagRep.13.5.182

1 Executive Summary

Longfei Shangguan
Xia Zhou
Marco Zimmerling

The past few years have witnessed progressive deployments of IoT devices in both personal and public space to facilitate smarter living. For instance, home gateways such as Amazon Echo can now pick up human speech, understand the semantics, and further loop in machines, appliances, and many others to react. Likewise, mobile IoT devices such as drones and robots are deployed for urban modeling, express delivery, and industrial inspection. Although the functional innovations builds upon existing IoT architecture are going full steam ahead, the newly emerging scenarios such as Internet of ocean and implantable things often come with limited power budget, form factor, and computation resources, which fundamentally challenge the conventional IoT architecture, communication primitives, and sensing capabilities.

* Editor / Organizer

Table 1 Agenda.

	Day One	
09:00 — 09:15	Opening remarks	Longfei Shangguan, Xia Zhou, Marco Zimmerling
09:15 — 10:15	Keynote speech	Victor Bahl (Microsoft)
10:15 — 10:45	Break	
10:45 — 12:00	Speak-and-spark I	
12:00 — 13:30	Lunch	
13:30 — 15:00	Speak and spark II	
15:00 — 15:40	Coffee Break	
15:40 — 16:20	Invited Talk I (Prof. Olga Saukh)	Embedded Intelligence: The Way Forward
16:20 — 17:45	Speak and spark III	
	Day Two	
09:00 — 09:45	Invited Talk II (Prof. Alberto Quattrini Li)	Current Snapshot and Opportunities on Underwater Sensing and Communication
09:45 — 10:30	Invited Talk III (Dr. Justin Chan)	Intelligent Mobile Systems for Equitable Healthcare
10:45 — 12:00	Breakout session I	
12:15 — 13:30	Lunch	
13:30 — 14:15	Invited Talk IV (Prof. Marco Zuniga)	Sunlight for Wireless Communication
14:15 — 15:00	Invited Talk V (Prof. Baccelli & Prof. Hahm)	IETF Protocols & Embedded Software for Low-Power IoT using RIOT
15:00 — 16:00	Coffee Break	
16:00 — 17:15	Breakout session II	
17:15 — 18:00	Sync on technical report writing	
	Day Three	
09:00 — 10:00	Plenary discussion I: underwater & medical & airborne	
10:00 — 11:00	Break	
11:00 — 12:00	Plenary discussion II: energy-efficiency & AI & sensing and communication	
12:00 — 12:15	Closing remarks	Longfei Shangguan, Xia Zhou, Marco Zimmerling

The goal of this Dagstuhl Seminar was to bring together researchers from both academia and industry from around the world to *i*) review the existing IoT research directions from radical perspectives; *ii*) figure out fundamental challenges in modern IoT application scenarios that may then be investigated in joint research projects; and *iii*) discuss new types of hardware architecture, network stack, and communication primitives for these emerging IoT scenarios. An essential part of this seminar was to bridge the gap between the research momentum from academia and the practical needs from the industry and to actively engage a dialog between different communities. As IoT by nature is a multidisciplinary area involving innovations in hardware and architecture, algorithm and protocol, as well as applications, we sought researchers from different domains that are open to different research perspectives and welcome research of a different nature.

Seminar Schedule

The seminar commenced with a keynote speech by Dr. Victor Bahl, a Technical Fellow at Microsoft. This was followed by a pair of speak-and-spark sessions, during which each attendee delivered a brief presentation on their research, facilitating mutual familiarization of everyone's areas of interest. The inaugural day culminated with an invited talk given by Dr. Olga Saukh. Moving to the subsequent day, the program featured four invited talks along with two break-out sessions. During these break-out sessions, attendees were organized into smaller groups, each focusing on a specific IoT research theme. The final day was dedicated to wrapping up the event and engaging in discussions. Leaders of each break-out group presented summary reports of their respective discussions, effectively concluding the seminar. The detailed agenda is listed in Table 1.

Outcomes

The seminar centered around six comprehensive research topics within the realm of the Internet of Things (IoT). These topics encompassed a diverse spectrum, including underwater communication, implantable IoT devices, airborne communication systems, integrated com-

Figure 1 Speak and spark session I.

Figure 2 Speak and spark session II.

munication and sensing technologies, the intersection of artificial intelligence and IoT, and the optimization of energy efficiency in IoT applications. Through in-depth presentations and break-out discussions on these subjects, participants gained valuable insights into the cutting-edge advancements, challenges, and opportunities shaping the future landscape of IoT technology. The seminar fostered a collaborative environment where participants exchanged ideas and discussed potential research collaborations. Overall, the seminar was considered a success, building the next cornerstone for the wireless Internet of Things.

Furthermore, the extensive research discussions held during the break-out sessions have ignited a multitude of captivating concepts for the upcoming wave of IoT applications. This enthusiasm culminated in the creation of four technical reports, each focusing on distinct research realms:

- Energy Efficiency of IoT (§4.1)
- Integrated Communication and Sensing (§4.2)
- Advancements in Medical IoT (§4.3)
- Airborne Internet of Things (§4.4)
- Impact of AI on IoT (§4.5)

These reports delve into the challenges faced and chart potential trajectories for future progress in these specified areas.

2 Table of Contents

3 Overview of Talks

3.1 Resilient and Secure IoT

Arash Asadi (TU Darmstadt, DE)

Joint work of Arash Asadi, Andrea Ortiz
Main reference Amir Ashtari Gargari, Andrea Ortiz, Matteo Pagin, Anja Klein, Matthias Hollick, Michele Zorzi, Arash Asadi: "Safehaul: Risk-Averse Learning for Reliable mmWave Self-Backhauling in 6G Networks", CoRR, Vol. abs/2301.03201, 2023.
URL https://doi.org//10.48550/arXiv.2301.03201

In this talk, I will delve into the endeavors of our research team aimed at enhancing mmWave self-backhauling. Our focus has been on addressing challenges related to beamforming and route selection within self-backhauled networks. We've explored solutions using reinforcement learning as well as traditional optimization methods. Additionally, I will elaborate on our strategies for countering adversarial WiFi sensing and mitigating user tracking through the application of radio fingerprinting techniques. Furthermore, I will provide a concise overview of our advancements in Reflective Intelligent Surfaces (RISs) development, encompassing both sub-6GHz implementations utilizing RF-switches and 60GHz designs employing Liquid Crystals.

3.2 IETF Protocols & Embedded Software for Low-Power IoT using RIOT

Emmanuel Baccelli (FU Berlin, DE) and Oliver Hahm (Frankfurt University of Applied Sciences, DE)

The Internet of Things (IoT) is a term generally used for a (too) wide variety of hardware, platforms and protocols. In this talk we focus less on cloud/edge IoT aspects and hardware based on CPU/GPU, and more on things operated "beyond" edge computing: we focus on ultra-low power hardware, software and protocols running on microcontrollers. This is a category of devices the IoT relies heavily upon. To be more concrete, RFC7228 distinguishes several classes of such devices, for example Class 1 gathers devices providing memory budgets of 10 kB RAM and 100 kB flash. With such a target in mind, we quickly overview the key aspects of relevant network protocols (distributed algorithms, standard specifications, interoperable implementations) and we flip through the list of relevant working groups that are currently active at the Internet Engineering Task Force (IETF/IRTF) standardizing IP protocol adaptations for ultra-low power IoT. In parallel, we overview the embedded software platform developments that take place, in order to provide open source interoperable implementations of these open standards. We overview developments based on RIOT, a small-footprint operating system that runs on a wide variety of microcontroller boards and aggregates various libraries and network stacks.

3.3 The Inevitable Unification of the Cloud and Telecommunications Infrastructures – Science, Technologies, Strategy and Opportunities

Victor Bahl (Microsoft Corporation – Redmond, US)

We are at the beginning of an unprecedented opportunity for information technology startups and the established cloud industry to become a part of the next generation telecommunications infrastructure. Together, we can radically change it through softwarization, AI and edge computing. I will describe the scientific advances and business needs for bringing us to where we are today and then cast an eye to the future in sharing with the audience a vision for where things are going with telecommunications, including key enablers and potential surprises on the horizon. This will set the context for describing the opportunity ahead for the engineering and research community in the next several years, and beyond, as we stay at the forefront of the modernization that will enable ubiquitous computing via telecommunication networks propelled by innovations in AI and edge computing. I will next move into near-term strategy with Microsoft standing up of a new business division called Azure for Operators (AFO). I will describe the vision that led to the creation of AFO, its mission, and products, and the significant technical and scientific challenges, which we are pursuing. When overcome they will lead to the inevitable convergence of two massive industries that will open up a new decade of opportunities for universities, research institutes, established companies, and startups.

3.4 Software-Defined Wireless Communication Systems

Bastian Bloessl (TU Darmstadt, DE)

Future wireless communication systems will be increasingly software-defined, with AI-native components to optimize performance. With current Software Defined Radios, there are powerful general-purpose hardware platforms that could drive such systems. What we are missing are software frameworks that can make efficient use of the heterogeneous compute resources available on those platforms. In this talk, I will briefly discuss my research effort on filling this gap by designing and implementing novel real-time signal processing systems that provide full control of the data flow through custom schedulers and native support for hardware accelerators like FPGAs and GPUs. This can serve as the enabler for a new generation of wireless networks that use AI not just for multi-objective optimizations but as the base for a new system design that goes beyond traditional protocol structures and complex error handling to create more robust and more resilient systems.

3.5 Intelligent Mobile Systems for Equitable Healthcare

Justin Chan (University of Washington – Seattle, US)

Access to even basic medical resources is greatly influenced by factors like an individual's birth country and zip code. In this talk, I will present my work on designing intelligent mobile systems for equitable healthcare. I will showcase three systems that are not only interesting from a computational standpoint but are also having real-world medical impact. The first system can detect ear infections using only a smartphone and a paper cone. The second system enables low-cost newborn hearing screening using inexpensive earphones. Lastly, I will present an ambient sensing system that employs smart devices to detect emergent and life-threatening medical events such as cardiac arrest. Through these examples, I will demonstrate how new computational and sensing techniques that generalize across hardware and work in real-world environments can help to address pressing societal problems.

3.6 Low-Power Internet-of-Things on Earth and Space

Akshay Gadre (University of Washington – Seattle, US)

As the number of connected devices increases, the stress on the wireless bandwidth available keeps increasing. Researchers across the area of wireless systems are building novel communication and sensing systems to overcome these limitations by addressing deployment domain-specific problems such as, Space-Based Internet, Low-Power Low-Bandwidth IoT Sensors for Agriculture, and Low-cost Sensing Solutions for the supply chain. Yet, many of these research prototypes are considered as research artifacts which provide basic connectivity and sensing, yet lacking much of the improved utility (as typically mentioned in many research contributions).

This talk will focus on the important wireless problems at the frontier of these domains to improve end-user acceptance and utilization of these impressive technologies. This talk will present several important questions that a typical user faces today and how academics and industry researchers can collaborate on these problems to alleviate these issues. Finally, this talk will provide the base for interesting discussion throughout the day of meetings with researchers on brainstorming these ideas and building collaborations.

3.7 Energy Efficient Sensing for IoT

Junfeng Guan (EPFL Lausanne, CH)

Main reference Junfeng Guan, Jitian Zhang, Ruochen Lu, Hyungjoo Seo, Jin Zhou, Songbin Gong, Haitham Hassanieh: "Efficient Wideband Spectrum Sensing Using MEMS Acoustic Resonators", in Proc. of the 18th USENIX Symposium on Networked Systems Design and Implementation (NSDI 21), pp. 809–825, USENIX Association, 2021.
URL https://www.usenix.org/conference/nsdi21/presentation/guan

I will explore the integration of Micro-Electro-Mechanical-Systems (MEMS) filters with innovative sparse recovery algorithms, enabling energy-efficient IoT devices to effectively capture wideband spectra. I will showcase the practical implementation of this approach in two distinct applications. Firstly, I will introduce an energy-conscious spectrum sensing approach that empowers low-power IoT devices to detect wideband channel occupancies. This capability enables these devices to opportunistically utilize unoccupied channels, facilitating dynamic spectrum sharing. Secondly, I will unveil a precise self-localization technique for IoT devices. This method harnesses existing ambient 5G communication signals without necessitating any alterations to the 5G infrastructure. This innovation offers accurate self-localization capabilities without requiring modifications to the underlying infrastructure.

3.8 UAV Systems and Networks

Karin Anna Hummel (JKU Linz, AT)

We aim to develop efficient mobile networked systems that include machines, vehicles, and humans. Hereby we propose novel architectures of such networked systems, analyze the performance, and experimentally evaluate the overall capabilities of the wireless network and often context-sensitive algorithms for the wireless link or overall networked system, including predictive approaches. Two recent application scenarios for which we set-up small wireless networked testbeds to work on novel solutions target systems for human-drone teams in indoor scenarios, where the drone's camera is the primary sensor, used to navigate and interact with the human and the environment. We find that a distributed control architecture leveraging direct Wi-Fi links is feasible and beneficial for such systems having an AI or machine learning component, yet common Wi-Fi cannot provide time guarantees which would be necessary for many real-time applications. However, this architecture enables the successful deployment of sophisticated applications such as drone learning from human emotions, which is one of our current research projects. In this talk, I will briefly introduce some of our current research results on these research threads.

3.9 Wireless Sensing and Interactions

Tianxing Li (Michigan State University – East Lansing, US)

In this talk, I will present three distinct advancements in the realm of sensing and interaction technologies. Firstly, I will introduce an innovative approach to inaudible attacks on smart earbuds, overcoming limitations of existing methods by leveraging both direct and reflective paths for attacking signals, thereby enhancing signal-to-noise ratios. Secondly, I will present a 3D Hand Posture Reconstruction system utilizing 2D Rolling Fingertips, which capitalizes on active optical labeling and exploits rolling shutter effects in smartphone cameras to achieve improved 3D tracking for complex scenarios like underwater environments and specialized applications like virtual writing. Lastly, I will present a low-power system for precise face touch detection, integrating wrist inertial and novel finger vibration sensors with a cascading classification model for efficient gesture filtering. This system achieves a high F-1 score of 93.5% for detecting face touch events, while maintaining minimal power consumption, making it suitable for prolonged usage with battery-powered devices. Each advancement is supported by practical prototypes and evaluations, showcasing their potential across a range of real-world scenarios.

3.10 Toward a Battery-less Internet of Things

Andrea Maioli (Polytechnic University of Milan, IT)

Joint work of Andrea Maioli, Luca Mottola, Mikhail Afanasov, Naveed Bhatti, Dennis Campagna, Giacomo Caslini, Fabio Massimo Centonze, Koustabh Dolui, Erica Barone, Muhammad Hamad Alizai, Junaid Haroon Siddiqui
Main reference Andrea Maioli, Luca Mottola: "ALFRED: Virtual Memory for Intermittent Computing", in Proc. of the SenSys '21: The 19th ACM Conference on Embedded Networked Sensor Systems, Coimbra, Portugal, November 15 – 17, 2021, pp. 261–273, ACM, 2021.
URL https://doi.org//10.1145/3485730.3485949

Battery-less devices revolutionize conventional battery-dependent setups by integrating ambient energy harvesting technologies, thereby enabling a sustainable Internet of Things (IoT) paradigm with minimal maintenance requirements. However, the intermittent nature of ambient energy sources presents challenges. These devices frequently encounter energy shortages, causing intermittent computation where active phases alternate with inactive ones due to waiting for adequate energy. Given their limited energy reservoir, optimizing battery-less device operation is crucial to maximize computational output from ambient energy. This presentation introduces our research effort addressing these concerns. We begin with the introduction of ALFRED, a virtual memory abstraction and compilation pipeline tailored for mixed-volatile platforms. ALFRED autonomously identifies optimal program state mappings between volatile and non-volatile memory, enhancing efficiency. Our focus also extends to ensuring consistent efficiency in battery-less devices. We've developed a system design to efficiently regulate supply voltage and clock frequency, even in hardware-limited devices lacking dynamic voltage and frequency scaling support. Implementing two hardware/software co-designs capturing these aspects, our approaches reduce battery-less device energy consumption by up to 170% and significantly decrease workload completion time, enhancing their overall performance. Through these advancements, we contribute to the realization of robust and optimized IoT systems powered by ambient energy sources.

3.11 Space-IoT

RangaRao Venkatesha Prasad (TU Delft, NL)

Joint work of RangaRao Venkatesha Prasad, Sujay Narayana

The Internet of Things (IoT) is making a great impact on our lives and changing the way we interact with others, the environment, and even machines. Now there are two considerations for IoT space. (a) The IoT devices deployed on the ground need a reliable connectivity backbone. (b) Space technology is seeing unprecedented growth since space subsystems are becoming smarter, more reliable wireless links, and further miniaturized – all these without the need for radiation-hardened components in Low Earth Orbit (LEO). The above two aspects lead us to a new domain of IoT called the Space Internet of Things (Space-IoT). In recent times, there has been a huge interest and investment in Space-IoT. Space can be a suitable platform to solve many of the current and future problems in IoT, and the possible solutions are yet to be explored in depth in the space environment. In this talk, I will discuss my group's research effort in this research domain.

3.12 Exploring the Aquatic World with an Internet of Underwater Robots and Sensors: Current Snapshot and Opportunities on Underwater Sensing and Communication

Alberto Quattrini Li (Dartmouth College – Hanover, US)

The aquatic world, with more than 70% of the Earth covered by water, plays a critical role in our society and the economy, as recognized by many organizations including the United Nations and the World Wildlife Fund, which estimated the economy value related to the aquatic world – also called Blue Economy – to be at least US$24 trillion. Yet, currently, more than 80% of the ocean is unmapped. This talk delves into the challenges and opportunities of deploying an Internet of Underwater Robots and Sensors (or more generally, Things) to explore the underwater world and support the Blue Economy, particularly focusing on underwater sensing and communication. Supported by lessons learned from actual field experiments, I will present first the general requirements for such underwater robots and sensors to operate in such environments; second, the technologies used in the current deployments, highlighting gaps and open problems; and finally, current research that shows research opportunities in the space of Ocean Internet of Things to achieve the long-term goal of ocean IoT, to support large scale aquatic applications, such as environmental monitoring or archaeological exploration.

3.13 Embedded Intelligence: A Way Forward

Olga Saukh (TU Graz, AT)

AI is utilized in various IoT applications, ranging from environmental monitoring to home automation and production processes. However, the most concerning applications of AI are in safety-critical systems, such as transportation, medicine, and control, because incorrect use of AI can have devastating consequences. A significant amount of current research effort is focused on increasing the robustness of deep models or implementing safe and low-resource post-deployment domain adaptation techniques to manage domain shifts. In this talk, I will provide a concise overview of the state-of-the-art in the field and present the latest research results from my team. We discover essential properties of the loss landscape geometry and leverage these to improve AI safety across a broad range of applications, particularly in scenarios where resource constraints are at stake.

3.14 Low Latency Data Downlink for Low Earth Orbit Satellites

Deepak Vasisht (University of Illinois – Urbana-Champaign, US)

Joint work of Deepak Vasisht, Jayanth Shenoy, Ranveer Chandra
Main reference Deepak Vasisht, Jayanth Shenoy, Ranveer Chandra: "L2D2: low latency distributed downlink for LEO satellites", in Proc. of the ACM SIGCOMM 2021 Conference, Virtual Event, USA, August 23-27, 2021, pp. 151–164, ACM, 2021.
URL https://doi.org//10.1145/3452296.3472932

Large constellations of Low Earth Orbit satellites promise to provide near real-time high-resolution Earth imagery. Yet, getting this large amount of data back to Earth is challenging because of their low orbits and fast motion through space. Centralized architectures with few multi-million dollar ground stations incur large hour- level data download latency and are hard to scale. We propose a geographically distributed ground station design that uses low-cost commodity hardware to offer low latency robust downlink. We also discuss intersections of compute and networks in this emerging context. We believe our work engineers a paradigm shift similar to the one from super-computers to data centers, where software architectures can extract high performance from commodity hardware.

3.15 Towards Ubiquitous Mobile Connectivity

Chenren Xu (Peking University, CN)

Mobility is the essential norm in human society. As today's wireless and mobile networking technology already brings 99% usable connectivity between our personal devices and (edge) cloud services, the remaining fragmented 1% connectivity scenarios, including but not limited to extremely lower power, high mobility, and massive access, which still face domain-specific networking challenges. More importantly, these minority cases are likely to turn over into the majority tomorrow as our global society is consciously evolving in the direction of energy and production efficiency improvement. This talk will introduce our recent efforts towards the vision of "Ubiquitous Mobile Connectivity". Specifically, we will present the design, implementation and deployment experience of our mobile RFID, VILD and multipath networking system for improving scalability, availability and reliability in logistics, Vehicular-to-X and high-speed railway networks.

3.16 The Small Data Problem in Understanding Older Adult Mobility

Rong Zheng (McMaster University – Hamilton, CA)

Joint work of Rong Zheng, Yujiao Hao, Boyu Wang

The proportion of older adults (aged 65 and older) worldwide has been increasing steadily over the past 40 years. Among older adults, mobility is a crucial indicator of functional status, and a predictor of quality of life and longevity; hence, it is often called the sixth vital sign. Mobility encompasses not only the physical activities of older adults, and the performance of specific maneuvers such as sit-to-stand, walking or climbing stairs, but also participation in society (e.g., the ability to drive, accessibility to public transportation). Recent advancements in mobile technologies, artificial intelligence, embedded and sensing devices create exciting opportunities to address mobility challenges faced by the aging population. However, the shortage of quality, labeled, free-living data from target populations remains to be a main barrier in developing and applying effective solutions for continuous monitoring, analysis and assessment.

My research work aims to mitigate the data scarcity problem in characterizing old adult activities and answer "what", "when" and "how well" questions regarding their functional mobility. We take a multi-pronged approach. Directly addressing data availability, we developed ChromoSim, a cross-mobility IMU data synthesizer, CROMOSim, that can transform abundant data from vision, motion capture systems to IMU data. At the modelling level, we devised efficient invariant feature learning framework that can easily adapt to novel subjects or devices and can handle noisy labeled data collected in the wild.

3.17 Sunlight for Wireless Communication

Marco Antonio Zúñiga Zamalloa (TU Delft, NL)

Joint work of Seyed Keyarash Ghiasi, Marco Antonio Zúñiga Zamalloa, Koen Langendoen
Main reference Seyed Keyarash Ghiasi, Marco Antonio Zúñiga Zamalloa, Koen Langendoen: "A principled design for passive light communication", in Proc. of the ACM MobiCom '21: The 27th Annual International Conference on Mobile Computing and Networking, New Orleans, Louisiana, USA, October 25-29, 2021, pp. 121–133, ACM, 2021.
URL https://doi.org//10.1145/3447993.3448629

We, humans, already use 50% more energy moving information than moving airplanes around the world. Communication is central to our societies but it is taking a toll on the earth. We want to use a free, abundant, and natural resource for wireless communication: sunlight. Similar to the way you can use a mirror to communicate by reflecting light, our aim is to

Figure 3 Plenary report and discussion I.

Figure 4 Plenary report and discussion II.

exploit optical devices that can change their reflective and transmissive properties to send information, but without you noticing any flicker. In this manner, objects will be able to talk to each other using daylight, an eco-friendly solution. In this talk, I will present some of our recent work in this area and its potential applications.

4 Breakout Session Report

4.1 Energy Efficiency of IoT

Tianxing Li (Michigan State University – East Lansing, US), Andrea Maioli (Politecnico di Milano – Milano, IT)

4.1.1 Relevance of energy-efficiency improvements

Better energy efficiency leads to performance improvement. However, such performance improvement may be irrelevant, depending on the application scenario and the device's lifetime. For example, in agriculture applications where deployments last for an entire season, a couple of days of improvement in the battery life is not significant enough. Conversely, in other application scenarios, such as battery-less devices deployments that use harvested energy, a slight improvement (i.e., 10%) in energy efficiency may lead to a lower number of energy failures or to energy/power neutrality (i.e., the system consumes the same amount of harvested energy, enabling perpetual and unattended operations). Therefore, the relevance of performance increase due to higher energy efficiency is application-specific.

4.1.2 Energy efficiency improvements target

IoT devices usually sense, compute, interact with the environment, and communicate sampled data. These four types of actions have different energy consumptions, which change depending on the IoT device and the application scenario. Unthinkingly improving energy efficiency may not result in a performance increase. For example, improving the energy efficiency of computation in underwater devices and drones does not lead to a significant performance increase, as their energy consumption is primarily due to movements. Similarly, devices that periodically sense the environment and enter deep sleep after each measurement consume the most energy while inactive. In such a scenario, improving the energy efficiency of sensing, processing, or communication leads to a lower performance improvement than improving deep-sleep energy efficiency. For these reasons, researchers should focus their resources on improving the energy efficiency of the most energy-hungry components of their applications.

4.1.3 Design space exploration

There are multiple platforms, techniques, and components to design IoT devices. These elements affect devices' energy efficiency. However, as we argued before, the most efficient components/techniques (i.e., custom designing an SoC) does not necessarily translate to the most efficient choice due to higher costs, higher design time, or irrelevant performance improvements. Therefore, when designing IoT devices, we must consider every trade-off that comes with using specific hardware and software components and select the one that suits the application requirements.

4.1.4 Energy harvesting

Ambient energy harvesting unlocks new scenarios where devices use the energy available in the environment to operate and recharge their batteries, extending battery life and reducing maintenance efforts. Energy harvesting can be used to improve devices' energy efficiency. During the seminar, we identified two complementary options. First, devices can harvest energy produced during normal operations. For example, we need to sense the airflow to stabilize drone's movement. In this case, we can harvest the energy from the movement of the drones to build a battery-less airflow sensing system. Second, devices can correlate sensing operations with energy harvesting sources. For example, we can harvest solar energy from ambient light intensity sensing to enable a battery-less gesture recognition system. In general, improving energy harvesting techniques was considered the "to-go" option during the seminar, with a particular interest in recovering the energy wasted due to normal operations (e.g., harvesting wind energy from drone movements).

4.1.5 Dynamic/runtime parameter adjustments

The increase in design complexity and the possibility of energy harvesting unlocks new optimization techniques. Devices should adapt their duty cycles and their application workload depending on available energy and the power source they are currently using (e.g., ambient energy or battery). Moreover, applications should dynamically adapt their parameters to always operate in the most efficient setting. Existing systems vary device operating settings, such as voltage and frequency. However, devices can also adapt other parameters, such as their sampling rate, the computation resolution, and the number of algorithm inputs (e.g., machine learning models). This requires designing new techniques to identify the parameters to tune and new energy-aware hardware/software designs that provide the required capabilities, such as power source recognition and real-time device energy consumption measurement.

4.1.6 Edge/local computation

The advances in AI applications increase the computational complexity of the processing executed on IoT devices, with a consequent increase in energy consumption and resource utilization. Local device computation may no longer be beneficial or possible due to tight energy budgets or insufficient resources (e.g., ML models may require more memory than the one available). Therefore, energy consumption due to data communication and sensing may no longer dominate. Data offloading may become an option on highly resource-constrained IoT devices, despite communication costs that usually represent the highest energy consumption of applications. Therefore, system designers should investigate the trade-offs between local and edge computation. In both these cases, data compression techniques may improve devices' energy consumption, as they can decrease the size of communication data or the size of ML models, potentially leading to lower energy consumption. However, the challenge is whether we can design a general framework to optimize the compression parameters on the fly so that the system can adapt to complex and dynamic scenarios.

4.2 Integrated Communication and Sensing

Bastian Bloessl (TU Darmstadt, DE), RangaRao Venkatesha Prasad (TU Delft, NL)

4.2.1 The Role of Interference Management

Today's wireless technologies like LoRa and NB-IoT scale surprisingly well with the number of nodes. One of the reasons for this is the RF environment, which is in most cases relatively static, given the fact that gateways are deployed at exposed locations, making it less likely that significant proportions of the Fresnel zone are frequently obstructed. This raises the question, whether interference management is still of prime interest. We believe that the topic should not drift out of focus for two reasons: (1) Even though LoRa and NB-IoT are popular choices today, it is not clear what we will use in the future. The design space for wireless communication technologies is large. With better energy harvesting and more power-efficient nodes, it seems likely that we will see technologies with more capable nodes that can cater to higher traffic demands. In these cases, interference management will play a more important role again. (2) Future IoT networks will become three dimensional, integrating airborne gateways mounted on UAVs, balloons, or satellites. These exposed gateways will cover larger geographic areas and, therefore, a much higher number of nodes that compete for spectrum.

4.2.2 The Role of Intelligent Surfaces

Reconfigurable Intelligent Surfaces (RISs) are a technology that could help to shape the RF environment to increase coverage and manage interference, as they allow steering reflections on a surface to guide signal propagation. Apart from diversity gains from reflected signals, this is particularly interesting for NLoS scenarios, where RISs can increase coverage by steering reflected beams to a target. While this technology promises great advantages, it is, at the moment, mainly discussed for mmWave technologies and indoor scenarios, where the expected benefits are particularly large. Whether the adoption in IoT deployments is technologically feasible and economically attractive has yet to be seen. Both the lower frequency bands and the different deployment scenarios might be problematic: It is unclear how well RISs can be adapted for sub-GHz bands and whether these surfaces are still cost effective, especially considering the extent of typical IoT networks. Another angle towards RISs is their potential use to enhance the privacy of users. Controlling signal propagation, they cannot only influence where signals go but also where they do not go. A signal that an eavesdropper never receives or overhears presents certainly the best protection to the user's privacy. Thinking further into the future, it would be interesting to extend the idea of RISs from reflection to penetration, working on materials that can change their attenuation coefficient. This could either be used to increase indoor reception or shield RF inside a room or building to protect from eavesdroppers.

4.2.3 (Joint) Communication and Sensing

The possibilities and accuracy of RF sensing is impressive, especially considering that most works exploit signals from ubiquitous technologies like WLAN that are not specifically designed for sensing. The fact that it is possible to see through walls, detect people, emotions, and even vital signs shows the potential. Future technologies will likely be designed with sensing in mind and waveforms will be adapted accordingly. One of the main questions that arises in this context is whether there should be a waveform for joint communication and sensing or if there are dedicated waveforms for communication and sensing. The former will likely be a compromise between the two applications, while the latter provides more flexibility to support sporadic sensing or increasing the accuracy by dedicating more spectrum to sensing. Besides technologies with native support for sensing, the most disruptive change in the field results from advancements in machine learning algorithms. Yet, we often reuse existing models for RF sensing. Prime examples are image networks that are used to process RF spectrograms. We believe that the next step for sensing is the development of foundational models, dedicated to the application. A network that processes IQ samples, i.e., the time domain signal, does not require calculating the FFT and includes phase information. Another question is how we train and feed these networks. Today's sensors are not designed to provide the input for ML algorithms. An audio device might, for example, do preprocessing with a psychoacoustic model of humans, filtering out information that might be relevant for ML. Or videos and pictures will be preprocessed to make them look nicer. This goes as far as replacing objects with reference pictures that were available during training (e.g. Samsung smartphones replace low-quality photos of the moon with high-resolution pictures). Future sensors might, therefore, have an option to output raw data or do preprocessing specifically for ML algorithms.

4.2.4 Privacy Implications and Ethics

Today, inventing new technologies comes with more responsibilities. Given the potential scale of IoT applications, privacy should not be an afterthought. Instead, we have to consider from the start what would happen, if the technology is adopted by thousands, millions, or billions. A recent negative example are Apple AirTags and similar trackers from other vendors. The potential for misusing the technology was obvious, yet, it was released on a global scale. Just now, Apple and Google work together on a standard to make misuse like stalking harder. Also more and more funding agencies expect considerations with the impact on privacy, society, or the environment. As researchers, we have a responsibility to educate the general public about the privacy implications of technologies, help them to make informed decisions, and ideally provide tools to protect their personal information. Considering the increasing complexity and ubiquity of IoT applications, the next big breakthrough could be a standard or technology that enables users to stay in control of their data.

4.3 The Future of Medical IoT Research

Justin Chan (University of Washington – Seattle, US), Rong Zheng (McMaster University – Hamilton, CA)

4.3.1 Challenges and recommendations

Community datasets. There is a need for more high quality biomedical datasets like PhysioNet and CheXNet to enable the community to benchmark technical progress in a similar fashion to the ImageNet Large Scale Visual Recognition Challenge.

Below we outline several challenges for high-quality dataset collection and curation:

Dataset diversity. To ensure that an IoMT system can scale, it is essential that any datasets on which it relies on is diverse across at least three different dimensions: a) patient demographics – These populations can involve demographics related to age, race and socioeconomic income. The demographics in the dataset should be representative of the patient population in which the system would eventually be deployed in. b) deployment environments – The data should be collected from a mixture of both controlled and uncontrolled environments in the wild. Controlled environments can include labs and clinics, while uncontrolled environments can include home environments. c) device heterogeneity – The data should be collected across different hardware models (e.g. different models of smartphones) and a clear description of how the data was collected should be described.

Dataset imbalance. Often in biomedical datasets, it is significantly easier to obtain data from a control population than from populations with a particular disease. To overcome issues of dataset imbalance, we believe there are opportunities to develop generative AI systems or physical simulations to create synthetic datasets that complement data collected from the real world.

Data normalization and standardized study protocols. IoMT devices are diverse in their form factors and physical properties. Placements and means of attachment also affect the characteristics of data obtained. For datasets to be exchangeable, ideally, standardized study protocols should be followed and some form of normalization need to be done. Minimally, the measurement procedure and detailed device specification should be provided along with the collected data.

Data science. Below we describe several areas that should be considered when analyzing biomedical datasets:

Clinical implications of system errors. Different medical applications can have different evaluation metrics. Even for false positive and false negative rates, what is considered acceptable is heavily dependent on the medical context. Analysis on biomedical datasets should describe the clinical implications of the metrics adopted, as well as how the error rates compare to standards set by regulatory and standards such as the FDA, ISO, or the performance obtained by similar medical technologies.

Subgroup analysis. Analysis of biomedical data should provide subgroup analysis to evaluate system performance across different dimensions including patient demographics, deployment environments, and system configurations. This will provide a more detailed picture of how the system performs instead of relying solely on aggregate performance numbers.

Explainable models. The development of explainable biomedical models for IoMT can help increase confidence in the system's decisions and pave the way to adoption. Specifically, explainable models can help to verify that the system is not making a decision based on artifacts in the data and is instead relying on clinically sound features.

Clinical studies

Establishing partnerships. Clinical studies often require access to clinical sites and collaborations with interested clinicians, which may not be easily accessible for everyone. We believe that there are opportunities beyond traditional healthcare sites where subjects can be recruited: 1) Partnering with public health authorities can be another path to collecting data for IoMT systems 2) Collaborations with private wearable or medical devices companies can enable collection of large-scale datasets 3) NGOs and global health organizations with existing clinical connections can enable clinical studies particularly in low and middle income countries 4) Crowdsourcing platforms can be a way to efficiently collect data from a diversity of environments and smart devices.

4.3.2 Opportunities

Personalized medicine. IoMT systems can enable longitudinal monitoring of an individual's health conditions, and adapt its sensing or detection capabilities to the individual through online or lifelong machine learning with private data. For example, there are opportunities for smart speakers and voice assistants to track health biomarkers to predict COPD relapses, dementia and Alzheimers.

Tracking of rare medical events. IoMT systems are able to capture significantly more data outside clinical settings especially in the home and can potentially track difficult to predict events like cardiac arrests and seizures. This can enable pre-screening and early diagnosis of medical disorders that would otherwise be challenging to track.

Remote diagnostics for telemedicine. With the increased adoption of telemedicine, there is an opportunity to create IoMT systems that can be performed remotely in conjunction with a physician or health care professional. These systems can involve repurposing the sensors on existing smart devices, or low-cost attachments that can be distributed at scale.

Ambient sensing systems. Ambient sensing systems can enable passive collection of medical data in domestic and public settings. Opportunities are abundant in the development of a) novel sensing methods leveraging chemical, acoustic, radio, or optical modalities in existing infrastructure and b) new privacy-preserving systems can enable large-scale epidemiological systems for tracking disease outbreaks.

Smart materials for wearable biofluid sensors. There are lots of opportunities to bring chemical testing out of the lab and create wearable and flexible chemical sensors that can continuously monitor biofluids like sweat, saliva, and tears to track vital signs and diseases.

Closed-loop systems. Beyond sensing, IoMT can be combined with actuators for medical intervention. Examples are, wearable sensors with auto-injectors that can detect and reverse events like opioid overdoses or anaphylaxis caused by allergic reactions; and prosthetics controlled by electromyography. The rapid progress in brain-computer interface (BCI) technology recently has opened up new possibilities in augmented or virtual reality applications with human-in-loop.

4.4 Airborne Internet-of-Things

Akshay Gadre (University of Washington – Seattle, US), Deepak Vasisht (University of Illinois – Urbana-Champaign, US)

Joint work of all members of the Airborne Internet-of-Things session

Internet-of-things is also unlocking the potential to augment devices not typically deployed statically as a sensor but instead a mobile device that can move around and sense different environments using the same sensor. One such direction of research exploration is understanding and augmenting the opportunity of airborne internet-of-things where devices with batteries/solar panels spend several seconds (insects), minutes (UAVs/drones/blimps), hours (airplanes) and even days (satellites) away from power.

There are tremendous opportunities that can be enabled by effectively capitalizing on the ability of these sensors to move around and provide connectivity. For instance, satellites and drones can act as data grabbers from terrestrial sensors with poor connectivity. Satellites and aircrafts can provide ocean coverage for various capabilities (localization, tracking, SOS, communication) that we traditionally take advantage of in well-connected environments. There has also been demonstration of using satellites for cellular backhauls for tracking global shipments by Telefonica. Further, there are new opportunities in exploring new capabilities such as edge computing and mesh networking in space. On the UAV side, enabling better autonomy for coordination of drone swarms for various sensing, transportation and entertainment applications remain important. Finally, planes can act as an intermediate zone of innovation between the drones which are battery constrained and the satellites which are communication constrained to assist the other two in enabling the same applications. It is also critical to not just build these services but also specializing them for commercial and environmental applications in real-world industries, such as mining, oil and gas, and agriculture. There are also upcoming deployments that cannot be classified into the above three categories but can enable new applications such as Insect IoT and Balloons IoT.

These opportunities can only be enabled by solving real world system challenges across the cyber-physical system stack from embedded and communication systems to better software applications that maximize the utilization of the limited resources at the sensors. One such challenge for satellites that is gaining increased exposure is the problem of reliability provided by satellite constellation-based internet. Indeed, a recent story of deployment in Ukraine highlighted this issue at a recent tutorial at MobiCom. Further, coordination of satellite communication directly to the sensors on the ground remains an open problem. Another important problem for satellite deployment is sustainability of deployment and debris reduction of existing satellites. In fact, recent efforts have taken place to use the 4th stage rockets lost in the orbit for satellite applications. Further, it was also described that most LEO satellite deployments deplane in orbit and get melted in the atmosphere. However, these are preliminary solutions and a more robust solution for sustainability at scale needs to be developed. Another important problem is standardization of communication protocols across the planet like cellular protocols where the constellations can effectively compete for the users without unnecessary augmenting overlapping infrastructure. On the UAV front, there are several traditional challenges of battery powered devices taken to the extreme – power, computation, storage, mobility and communication in both rural and urban environments.

As we move towards a more data-driven world, the role of AI becomes more inevitable. Especially with such mobile cyber-physical systems, it is critical to think about the guard rails and data-resilient technologies that remain safe for usage in human-coexistent environments. Further, parallel research in data sciences need to build tools and solutions that are trustworthy, reliable, and perhaps provable. An important aspect of that provable aspect will be the explainability of the model to describe what are the key features that had the most impact on the decision it took. As we learnt during a live talk by Olga Saukh, "the robustness of models to data drifts and sensors capability drifts will be critical to such sensors' and data-driven algorithms' ability to sustain equitable AI for cyber-physical systems." It will also be critical to build solutions that understand and evaluate the data for data biases and build automated tools to understand the biases of black-box AI models. Further, human-in-the-loop models should create more accessible interfaces for the human to explain the deviation in the decision-making process that it is introducing and the logic behind it needs to be learnt by AI algorithms on-the-line. Many of the more energy hungry AI models will need energy-elastic capabilities to continuously manage the energy-accuracy tradeoff during live applications. Finally the continuous verification of these AI systems and detection of maleficent intent is critical for a safer future.

4.5 Impact of AI on IoT

Olga Saukh (TU Graz, AT), Marco Antonio Zúñiga Zamalloa (TU Delft, NL)

The impact of AI on IoT research has been significant and transformative. Integration of ML models into IoT systems has led to numerous advancements across various domains, including medical, underwater and space IoT applications discussed at the seminar.

As increasingly more data is generated by IoT sensors, ML algorithms running on edge devices extract valuable insights, identify patterns, and make predictions based on the collected data. This facilitates the development of new diagnostic tools and real-time decision-making. The approach is particularly valuable in manufacturing, but also builds the core of early warning systems in medicine, precision agriculture and cattle farming.

AI-powered algorithms can help to optimize network management and communication protocols in IoT systems. This improves connectivity, enhances reliability, and reduces latency issues. AI can play a crucial role in enhancing the privacy and security of IoT systems. ML algorithms can analyze network traffic patterns and detect anomalies that indicate potential cyber threats. They can help to anonymize and encrypt sensitive IoT data, ensuring privacy protection.

Natural language interfaces and vision-based tools can enable seamless interactions between humans and IoT devices, creating more intuitive and user-friendly interfaces and experiences.

4.5.1 Challenges

Designing safe AI-based systems. AI-based systems in the IoT domain pose challenges such as distribution shift, where the real-world data encountered during deployment differs from the training data, robustness against adversarial attacks and noisy sensor data, and adaptation to dynamic environments. Addressing these challenges requires techniques such as continuous monitoring of data distributions, domain adaptation, adversarial training, robust optimization, and online learning to ensure the reliability, robustness, and adaptability of AI models in IoT systems.

Tackling scalability and resource constraints. Scaling AI-based IoT systems to handle a large number of connected devices and manage massive data streams can be challenging. AI algorithms may require significant computational resources, which can strain the limited resources of IoT devices. Optimizing AI algorithms for resource-constrained environments becomes necessary. Since integrating intelligence into IoT devices may become an unsolvable challenge, offloading parts of data processing to the cloud is a viable option, yet has to carefully take into account limited bandwidth and increased unbounded latencies.

Guaranteeing data quality and privacy. IoT systems generate massive amounts of data, but ensuring data quality and privacy can be challenging. AI models heavily rely on high-quality data for accurate results, and ensuring data integrity, security, and privacy protection becomes crucial. Addressing these challenges requires robust data governance and security measures on edge devices.

Enabling trustworthy AI-based tools. Explainability and interpretability of AI models in IoT devices pose a significant obstacle. Service providers and regulatory bodies require transparency to trust the decisions made by AI models. However, many advanced AI techniques are inherently complex and black-box in nature, making it challenging to provide understandable explanations for their outputs. The interpretability challenge is amplified in IoT devices where computational resources are at stake, limiting the feasibility of deploying complex interpretable models. Striking a balance between accuracy and interpretability is crucial in ensuring the safe and ethical deployment of AI.

Expanding skills and expertise. Developing AI-based IoT systems requires a mix of skills and expertise in AI, IoT, data engineering, and domain knowledge. Finding professionals with interdisciplinary skills is challenging. Promoting cross-domain collaboration becomes necessary.

4.5.2 Opportunities

Domain-specific AI solutions. The IoT community plays a crucial role in advancing domain-specific AI solutions by providing insights into challenges and requirements unique to various fields. By sharing anonymized and appropriately labeled data sets collected from diverse IoT sensors, the community can enable AI researchers to develop tailored models and benchmark their performance effectively. However, while data is crucial, there is often an imbalance in focus, with more emphasis on model development rather than data collection. To address this, a culture of sharing and reusing datasets needs to be cultivated, encouraging the IoT community to contribute to open data repositories and facilitating the exchange of valuable data for the advancement of AI in IoT applications. Data reuse is often difficult in the IoT community, since in many cases shared data tightly reflects peculiarities of a specific setup. Therefore sharing large rigorous and well-curated datasets for selected applications and updating these data over time can ultimately benefit the entire IoT ecosystem.

Sensors and AI-based edge systems. There is a need for a stronger hardware-model co-design to optimize sensors and enhance overall AI-based system performance. In the IoT domain, where multimodal sensing is prevalent, we call for exploring the appropriate modeling and representation learning techniques for different modalities. The current emphasis in AI research on images and text data should be expanded to encompass the multimodal nature of IoT data, which includes a wide variety of sensors. Additionally, the IoT community should question the applicability of conventional machine learning practices inherited from other domains. The latest results in sparsity suggest that the best performing sparsity methods for computer vision fail on large language models. It is essential for the IoT community to focus on identifying foundational models specific to IoT data that can capture the underlying dynamics and characteristics of the IoT environments.

Real-world use cases and knowledge sharing. Collaboration between the IoT community and AI researchers opens up opportunities for innovation and the development of tailored solutions that meet the specific demands of diverse IoT applications. The IoT community can share their real-world use cases and challenges with AI researchers, providing insights into specific domain requirements and practical scenarios. This sharing can take various forms, including data, generative models, simulations, deployment traces and lessons learned. By collaborating with AI researchers, the IoT community can help identify and address the challenges associated with IoT-related issues such as data silos, distribution shifts, domain adaptation and the need for continual learning. Furthermore, integrating simulators, game engines, and digital twins into IoT systems can enhance the development and testing of AI algorithms in realistic and open-world environments. The combination of physics and machine learning models also holds promise for advancing AI in the IoT domain, enabling more accurate and interpretable solutions.

Performance evaluation of AI-based systems. AI algorithms need to perform efficiently and effectively in IoT environments. The IoT community can actively participate in evaluating AI model compression techniques for their suitability for resource-constrained IoT devices and various hardware architectures, considering factors like energy consumption, computational requirements, and latency. This feedback can guide AI researchers in optimizing their algorithms and optimization techniques for practical settings. In addition, the verification of edge AI, running in parallel to current efforts, will emerge as a distinct field that addresses the unique requirements and considerations of AI algorithms deployed at the edge. The focus on performance evaluation and verification will contribute to the development of robust and efficient AI-based systems for IoT applications.

Ethical considerations. Deploying AI in IoT systems raises ethical concerns, such as privacy, security, bias, and transparency. The decision-making capabilities of AI algorithms should align with ethical and legal frameworks. Ensuring fairness, accountability, and transparency in AI-based IoT systems requires careful design, monitoring, and governance. The IoT community plays a crucial role in establishing standards and protocols for IoT devices and systems. Similarly, the AI community can contribute to defining standards for AI-enabled IoT devices, ensuring interoperability, data exchange, and security. Collaborative efforts can lead to the development of standardized interfaces and frameworks that facilitate the integration of AI and IoT technologies.

Participants

- Arash Asadi
TU Darmstadt, DE
- Emmanuel Baccelli
FU Berlin, DE
- Victor Bahl
Microsoft Corporation – Redmond, US
- Bastian Bloessl
TU Darmstadt, DE
- Justin Chan
University of Washington – Seattle, US
- Akshay Gadre
University of Washington – Seattle, US
- Junfeng Guan
EPFL Lausanne, CH
- Oliver Hahm
Frankfurt University of Applied Sciences, DE
- Richard Han
Macquarie University – Sydney, AU
- Karin Anna Hummel
Johannes Kepler Universität Linz, AT
- Tianxing Li
Michigan State University – East Lansing, US
- Andra Lutu
Telefónica – Madrid, ES
- Andrea Maioli
Polytechnic University of Milan, IT
- RangaRao Venkatesha Prasad
TU Delft, NL
- Alberto Quattrini Li
Dartmouth College – Hanover, US
- Olga Saukh
TU Graz, AT
- Longfei Shangguan
University of Pittsburgh, US
- Deepak Vasisht
University of Illinois – Urbana-Champaign, US
- Chenren Xu
Peking University, CN
- Rong Zheng
McMaster University – Hamilton, CA
- Xia Zhou
Columbia University – New York, US
- Marco Zimmerling
TU Darmstadt, DE
- Marco Antonio Zúñiga Zamalloa
TU Delft, NL