DAGSTUHL
REPORTS

Volume 11, Issue 5, June 2021

ISSN 2192-5283

Published online and open access by
Schloss Dagstuhl – Leibniz-Zentrum für Informatik GmbH, Dagstuhl Publishing, Saarbrücken/Wadern, Germany. Online available at
http://www.dagstuhl.de/dagpub/2192-5283

Publication date
December, 2021

Bibliographic information published by the Deutsche Nationalbibliothek
The Deutsche Nationalbibliothek lists this publication in the Deutsche Nationalbibliografie; detailed bibliographic data are available in the Internet at http://dnb.d-nb.de.

Digital Object Identifier: 10.4230/DagRep.11.5.i

Aims and Scope
The periodical *Dagstuhl Reports* documents the program and the results of Dagstuhl Seminars and Dagstuhl Perspectives Workshops.
In principal, for each Dagstuhl Seminar or Dagstuhl Perspectives Workshop a report is published that contains the following:
- an executive summary of the seminar program and the fundamental results,
- an overview of the talks given during the seminar (summarized as talk abstracts), and
- summaries from working groups (if applicable).

This basic framework can be extended by suitable contributions that are related to the program of the seminar, e. g. summaries from panel discussions or open problem sessions.

Contact
Schloss Dagstuhl – Leibniz-Zentrum für Informatik
Dagstuhl Reports, Editorial Office
Oktavie-Allee, 66687 Wadern, Germany
reports@dagstuhl.de
http://www.dagstuhl.de/dagrep

Report from Dagstuhl Seminar 21231

Transparency by Design

Edited by

Judy Kay[1], Tsvi Kuflik[2], and Michael Rovatsos[3]

1 **The University of Sydney, AU,** judy.kay@sydney.edu.au
2 **Haifa University, IL,** tsvikak@is.haifa.ac.il
3 **University of Edinburgh, GB,** michael.rovatsos@ed.ac.uk

Abstract

This report documents the program and outcomes of Dagstuhl Seminar 21231 on "Transparency by Design" held in June 2021. Despite extensive ongoing discussions surrounding fairness, accountability, and transparency in the context of ethical issues around AI systems that are having an increasing impact on society, the notion of transparency – closely linked to explainability and interpretability – has largely eluded systematic treatment within computer science to date.

The purpose of this Dagstuhl Seminar was to initiate a debate around theoretical foundations and practical methodologies around transparency in data-driven AI systems, with the overall aim of laying the foundations for a "transparency by design" framework – a framework for systems development methodology that integrates transparency in all stages of the software development process. Addressing this long-term challenge requires bringing together researchers from Artificial Intelligence, Human-Computer Interaction, and Software Engineering, as well as ethics specialists from the humanities and social sciences, which was a key objective for the four-day seminar conducted online.

Seminar June 6–10, 2021 – http://www.dagstuhl.de/21231
2012 ACM Subject Classification Computing methodologies → Artificial intelligence; Human-centered computing → Human computer interaction (HCI); Software and its engineering → Software development techniques
Keywords and phrases Artificial Intelligence, Dagstuhl Seminar, Ethics, Human-ComputerInteraction, Software Engineering, Transparency
Digital Object Identifier 10.4230/DagRep.11.5.1
Edited in cooperation with Bogina, Veronika, Orphanou, Kalia, Podoletz, Lena, Shulner-Tal, Avital

1 Executive Summary

Judy Kay (The University of Sydney, AU)
Tsvi Kuflik (Haifa University, IL)
Michael Rovatsos (University of Edinburgh, GB)

As AI technologies are witnessing impressive advances and becoming increasingly widely adopted in real-world domains, the debate around the ethical implications of AI has gained significant momentum over the last few years. Much of this debate has focused on fairness, accountability, transparency and ethics, giving rise to "Fairness, Accountability and Transparency" (FAT or FAccT) being commonly used to capture this complex of properties as key elements to ethical AI.

However, the notion of transparency – closely linked to terms like explainability, accountability, and interpretability – has not yet been given a holistic treatment within computer science. Despite the fact that it is a prerequisite to instilling trust in AI technologies, there is a gap in understanding around how to create systems with the required transparency, from demands on capturing their transparency requirements all the way through to concrete design and implementation methodologies. When it comes to, for example, demonstrating that a system is fair or accountable, we lack usable theoretical frameworks for transparency. More generally, there are no general practical methodologies for the design of transparent systems.

The purpose of this Dagstuhl Seminar was to initiate a debate around theoretical foundations and practical methodologies with the overall aim of laying the foundations for a "Transparency by Design" framework, i.e. a framework for systems development that integrates transparency in all stages of the software development process.

To address this challenge, we brought together researchers with expertise in Artificial Intelligence, Human-Computer Interaction, and Software Engineering, but also considered it essential to invite experts from the humanities, law and social sciences, which would bring an interdisciplinary dimension to the seminar to investigate the cognitive, social, and legal aspects of transparency.

As a consequence of the Covid-19 pandemic, the seminar had to be carried out in a virtual, online format. To accommodate the time zones of participants from different parts of the world, two three-hour sessions were scheduled each day, with participant groups of roughly equal size re-shuffled each day to provide every attendee with opportunities to interact with all other participants whenever time difference between their locations made this possible in principle. Each session consisted of plenary talks and discussion as well as work in small groups, with discussions and outcomes captured in shared documents that were edited jointly by the groups attending different sessions each day.

The seminar was planned to gradually progress from building a shared understanding of the problem space among participants on the first day, to mapping out the state of the art and identifying gaps in their respective areas of expertise on the second day and third day.

To do this, the groups identified questions that stakeholders in different domains may need to be able to answer in a transparent systems, where we relied on participants to choose domains they are familiar with and consider important. To identify the state of the art in these areas, the group sessions on the second and third days were devoted to mapping out the current practice and research, identifying gaps that need to be addressed.

The two sessions on each day considered these in terms of four aspects: data collection techniques, software development methodologies, AI techniques and user interfaces.

Finally, the last day was dedicated to consolidating the results towards creating a framework for designing transparent systems. This began with each of the parallel groups considering different aspects: Motivating why transparency is important; challenges posed by current algorithmic systems; transparency-enhancing technologies; a transparency by design methodology; and, finally, the road ahead.

The work that began with the small group discussions and summaries continued with follow up meetings to continue the work of each group. The organisers have led the work to integrate all of these into an ongoing effort after the seminar, aiming to create a future joint publication.

2 Table of Contents

3 Overview of Talks

3.1 The Limits of Transparency

Joanna J. Bryson (Hertie School of Governance – Berlin, DE)

Main reference Joanna J. Bryson: "The Artificial Intelligence of the Ethics of Artificial Intelligence", The Oxford Handbook of Ethics of AI, p. 1, Oxford University Press, 2020.
URL https://doi.org/10.1093/oxfordhb/9780190067397.013.1

Among many strong and positive suggestions in the 2020 EU whitepaper on AI was at least one repeated falsehood: that AI is necessarily opaque. AI is of necessity no more opaque than natural intelligence; in fact, digital artifacts can by choice be made to be far more transparent than nature. In the talk, I describe technological, sociological, and economic barriers to transparency, how these are affected by AI and the digital revolution, and what governance policies may be deployed to address them. Here in this extended abstract, I just talk about what transparency means, and what it is for.

My intent here is to provide the definitions most useful for interpreting the five OECD Principles of AI, since that's the soft law with the most international support, with 50 national governments (the OECD and the G20) now signed up to it.

Responsibility is the keystone. It is a property assigned by a society to individuals for their actions, including inactions where action was the expected norm. Actors with responsibility are technically termed 'moral agents' in philosophy. These vary by society. So for example, a family may hold a child or a dog responsible to know where an appropriate place is to pee. But a government will only hold legal persons responsible. In the case of a family, those will be the adult humans in the household.

Accountability is the capacity to assign responsibility to the correct agency. The purpose of accountability and responsibility are to maintain social order, that is, to maintain the society. Therefore responsibility is ordinarily assigned to those who can be held to account. The purpose of holding people to account is to persuade them and others like them to perform what a society considers to be their responsibilities for maintaining that society. Sometimes responsibility is determined to fall outside the control of any agency. In the USA for example, weather events are frequently termed "acts of God."

Transparency is the property of a system whereby it is possible to trace accountability and allocate responsibility. Where there is transparency, there does not need to be blind trust. Formally, trust is the expectation of good behaviour afforded (by a truster) to others (trustees) where the trustee's behaviour is actually unknown and where the trustee actually has autonomy with respect to the truster. Transparency may create a sensation of trust, but it renders the formal state of trust unnecessary. Transparency is something that can and should be designed into an intelligent system.

I postulate in my talk that the limits of transparency are problems like computational combinatorics, which means we can never know everything in detail; sociological problems like political polarisation and identity politics, which make people not likely to believe information even if it is in front of them; and mutually exclusive goals, so for example there can be no meeting of the minds when one mind is focussed only on wealth creation and another mind is focussed only on ethics red lines. But AI does not make any of these problems worse. The fact that AI systems contain complicated components is no more of a problem than organisations composed of humans with complicated brains. We only need transparency as to who is responsible for ensuring a system has been developed and operated in such a way that if it functions incorrectly or to malign purpose, we can know who caused it to do such a thing, even if that cause is failure to follow best practice or the release of an inadequately tested or understood system.

3.2 Fairness-Aware Recommender Systems

Robin Burke (University of Colorado – Boulder, US)

Recommender systems are machine learning systems that provide personalized results to users across a wide array of applications from social media to e-commerce to news to online dating. As fairness in machine learning has become a major sub-field of research in the past five years, recommender systems have also benefited from this emphasis. However, as these fairness-aware systems begin to be deployed, it becomes quite clear that we know very little about how users think about and interact with systems that take ethical stances, stances which might put them at odds with user interests and goals. The multi-sided nature of recommender systems also becomes clear when we consider fairness and this suggests the need for transparency toward providers of items being recommended. Fairness-aware recommendation therefore raises two key challenges: (1) how best to implement transparency for users who consume recommendations designed to be fair and (2) how to implement transparency for item providers for whom fairness may be important but where transparency may enable adversarial manipulation by such users.

3.3 Artificial Intelligence for Social Good: When Machines Learn Human-like Biases from Data

Aylin Caliskan (University of Washington – Seattle, US)

Developing machine learning methods theoretically grounded in implicit social cognition reveals that unsupervised machine learning captures associations, including human-like biases, objective facts, and historical information, from the hidden patterns in datasets. Machines that learn representations of language from corpora embed biases reflected in the statistical regularities of language. Similarly, image representations in computer vision contain biases due to stereotypical portrayals in vision datasets. On the one hand, principled methodologies for measuring associations in artificial intelligence provide a systematic approach to study society, language, vision, and learning. On the other hand, these methods reveal the potentially harmful biases in artificial intelligence applications built on general-purpose representations. As algorithms are accelerating consequential decision-making processes ranging from employment and university admissions to law enforcement and content moderation, open problems remain regarding the propagation and mitigation of biases in the expanding machine learning pipeline.

3.4 Towards Personalized Explainable AI

Cristina Conati (University of British Columbia – Vancouver, CA)

The AI community is increasingly interested in understanding how to build artifacts that are accepted and trusted by their us-ers in addition to performing useful tasks. It is undeniable that explainability can be an important factor for acceptance and trust. However, there is still limited understanding of the actual relationship between explainability, acceptance, and trust and which factors might impact this relationship. In particular, although existing research on Explainable AI (XAI) suggests that having AI systems explain their inner workings to their end users can help foster transparency, interpretability, and trust. there are also results suggesting that such explanations are not always wanted by or beneficial for all users. These results indicate that research in XAI needs to go beyond one-size-fits-all explanations and investigate AI systems that can personalize explanations of their behaviors to the user's specific needs.

There is general agreement that such needs may depend on context, e.g., the type of AI application and criticality of the targeted tasks, but there is also evidence that, given the same context, *user differences* play a role in defining when and how explanations may be useful and effective.

These results call for the need to investigate *personalized XAI*, namely how to create AI systems that understand to whom, when and how to deliver effective explanations of their actions and decisions.

AI-driven personalization has been an active field of research for several decades, spanning fields such as recommender systems, intelligent-tutoring systems, conversational agents, and affect-aware systems. To provide personalization, an AI system needs to have an adaptive loop in which it acquires a model of its user by inferring relevant user properties from available observations and decides how to personalize its behavior accordingly, to favor at best the goal of the interaction

We see explanations as yet another element of personalization in the adaptive loop, where the system ascertains if and how to justify its behavior to the user based on its best understanding of user properties specifically relevant to evaluate the need for explanation. What these relevant properties are is still largely unknown, hence, a key step toward personalized XAI is re-search to fill this gap.

Two general types of user properties have shown to be relevant for personalization: long-term traits that do not usually change over short periods of time (such as cognitive abilities and personality traits); and transient short-term states such as attention, interest and emotions.

We argue that, given a specific AI application, different types and forms of explanations may work best for different users, and even for the same user at different times, depending to some extent on both their *long-term traits* and *short-term states*. As such, our long-term goal is to develop personalized XAI tools that adapt dynamically to the user's needs by taking both these types of user factors into account.

In this talk, I focus on research investigating the impact of *long-term traits*, and how they may drive personalization. I present a general methodology to address these two questions, followed by an examples of how it was applied to gain insights on which long-term traits are relevant for personalizing explanations in an intelligent tutoring system (ITS). I discuss how to move forward from these insights, and present research paths that should be explored to make personalized XAI happen.

References

1 C. Conati, O. Barral, V. Putnam, and L. Rieger. Toward personalized XAI: A case study in intelligent tutoring systems, *Artificial Intelligence*, 298, 2021.

3.5 Increasing Transparency with Humans in the Loop

Gianluca Demartini (The University of Queensland – Brisbane, AU)

Bias appears in data collected from human annotators. It is then propagated into the Artificial Intelligence (AI) models trained with such labelled datasets. Bias in AI is then presented to end users who interact with the AI-powered system with their own bias and stereotypes. In such a setting, increasing the level of transparency could be an alternative to popular approaches aimed instead at removing bias from the system.

In this talk, I first present an example from our recent research of bias present in labelled datasets generated by human annotators in the context of crowdsourced judgements of information truthfulness. I then discuss how data-driven models of human annotator interaction behaviour may be leveraged to better understand the behavioural diversity present in a group of human annotators and the potential bias reflected in the labels generated by them. Finally, I discuss possible approaches to manage such bias in data and AI going beyond the classic aim of removing it.

References

1 La Barbera, D., Roitero, K., Demartini, G., Mizzaro, S., and Spina, D. (2020, April). Crowdsourcing Truthfulness: The Impact of Judgment Scale and Assessor Bias. In European Conference on Information Retrieval (pp. 207-214). Springer, Cham.

2 Han, L., Roitero, K., Gadiraju, U., Sarasua, C., Checco, A., Maddalena, E., and Demartini, G. (2019, January). All those wasted hours: On task abandonment in crowdsourcing. In Proceedings of the Twelfth ACM International Conference on Web Search and Data Mining (pp. 321-329).

3 JHan, L., Checco, A., Difallah, D., Demartini, G., and Sadiq, S. (2020, October). Modelling User Behavior Dynamics with Embeddings. In Proceedings of the 29th ACM International Conference on Information and Knowledge Management (pp. 445-454).

3.6 Transparency by design

Virginia Dignum (University of Umeå, SE)

In this talk, discuss the desirability and challenges of a design approach to transparency and proposed it to be complemented and extended by a "transparency in design" approach that focus on the processes, choices and stakeholders. Current work on transparency by design focuses on algorithmic transparency, namely on the data, results and functionalities of algorithms but often ignores the context in which algorithms are developed and used, and the power relations that influence decisions and requirements. Moreover, algorithmic transparency is not always possible nor desirable (IP, security, complexity, etc) so it needs to be complemented by methods that build trust, such as contracts and formal governance processes.

3.7 Transparency issues for the Use of AI and Analytics in Financial Services

Ansgar Koene (EY Global – London, GB)

Data analytics, risk modelling and prediction are common practices with a long tradition in the financial services. Despite great interest and a large number of pilot projects to explore the use of AI for enhancing the power of these computational practices, however, established financial institutions have been slow to put AI into operational use for core financial activities. Key factors that are holding back the deployment of AI are concerns about over the ability to comply with regulatory requirements regarding the explainability and robustness of financial models.

In this talk I review existing regulatory requirements and compliance considerations for the use of data analytics in the financial industry, which must be considered when evaluating the potential for the use of AI in this sector. Starting with an overview of levels of governance requirements that apply to the use of models in financial services and an associated algorithm risk tiering I provide some examples of typical model governance approaches for different types of regulated model classes. Based on these established model governance requirements I consider four dimensions of risk associated with AI: the data; the models and their implementation; the approach to modelling; and the (lack of) accumulated experience for these types of models. When considering the mitigation of these risks, I briefly present some model validation and monitoring considerations for conceptual soundness, data quality review, outcomes analysis, implementation controls and performance monitoring that are associated with business knowledge, data governance, cross validation, IT governance and model after-care, respectively.

Zooming in on model explain ability, I consider the familiar issue of trade-off between model interpretability and accuracy by focusing in on additional objectives that interpretability serves, such as: debugging and improvement of the model; trust and acceptance; regulatory compliance; ethics; safety and transferability; and discovery of unknown relationships in the data by being better able to interpret model outcomes. Particular attention is paid to the challenge of addressing the different explainability needs of various stakeholders, as highlighted in the Bank of England working paper No816 on "Machine learning explainability in finance" an application to default risk analysis.

I conclude with some findings of a survey on the ways in AI is being used by banks, or how they anticipate using it in the near future.

3.8 Engineering Traceability: A Lens Connecting Transparency Tools to Accountability Needs

Joshua A. Kroll (Naval Postgraduate School – Monterey, US)

Accountability is widely understood as a goal for well governed computer systems, and is a sought-after value in many governance contexts. But how can it be achieved? Many authors suggest it is enabled by transparency, though without a clear mechanism or requirements this, too, is challenging to achieve adequately. Recent work on standards for governable artificial intelligence systems offers a related principle: traceability. Traceability requires establishing not only how a system worked but how it was created and for what purpose, in a way that explains why a system has particular dynamics or behaviors. It connects records of how the system was constructed and what the system did mechanically to the broader goals of governance, in a way that highlights human understanding of that mechanical operation and the decision processes underlying it. We examine the ways that traceability links transparency demands to accountability needs, distill from these a set of requirements on software systems driven by the principle, and systematize the technologies available to meet those requirements. From our map of requirements to supporting tools, techniques, and procedures, we identify gaps and needs separating what traceability requires from the toolbox available for practitioners. This map reframes existing discussions around accountability and transparency, using the principle of traceability to show how, when, and why transparency can be deployed to serve accountability goals and thereby improve the normative fidelity of systems and their development processes.

3.9 Transparency and the Fourth AI Revolution

Loizos Michael (Open University of Cyprus – Nicosia, CY)

Facilitated by the desire to scientifically understand and replicate human intelligence in machines, the First AI Revolution had as a primary consequent the offloading of the cognitive burden of humans – reminiscent of the offloading of the physical burden of humans during the First Industrial Revolution – with humans retaining control as domain experts in their interaction with machines. This central role of humans was considerably diminished during the Second AI Revolution, where the advent of Deep Learning drove humans to the subsidiary role of "blue-collar" workers annotating data over the machine learning "assembly line" – echoing the primary consequent of the Second Industrial Revolution – and raised concerns on humans ceding too much control to AI.

The ongoing effort towards building transparent AI can, ultimately, be seen as a natural reaction and a potential remedy to these concerns. In this context, post-hoc transparency on why a system exhibits a certain behavior is not sufficient. Rather, transparency mechanisms should be built by design into an AI system to reveal why the latter's exhibited behavior is what it should be, according to any applicable legal, social, and ethical frameworks. The direction foreshadowed by the need for such transparency is that of an oncoming Fourth AI Revolution, facilitated – as with the Fourth Industrial Revolution – by an increased communication between humans and machines, and a resulting transition from machine automation to machine autonomy that is guided, monitored, and evaluated by humans, towards building long-term trust and offering reassurances that machines will respect the values that humans deem important.

In this talk we argue that this form of transparency can be achieved without abandoning the central role that machine learning has played in AI so far, but by extending the role of humans from that of data annotators to that of machine coaches. Much like how humans help each other acquire new knowledge by interactively providing feedback, the Machine Coaching paradigm that we put forward proceeds on the basis that humans and machines engage in a dialog on why a certain behavior was exhibited by a machine: the machine provides an explanation based on its current knowledge, and if the human finds the explanation lacking then the latter effectively provides back an improved or more appropriate explanation that the machine proceeds to integrate into its learned knowledge for future use. Feedback bilaterally exchanged between humans and machines in the context of Machine Coaching is, therefore, specific to the given situation, and is provided in-situ and in an agile and dialectical manner, ensuring that the process is cognitively light for the human coach, and that the knowledge acquired by the machine is – by design and provably, in a formal sense – acceptable to the human with whom the machine is interacting.

References

1 Loizos Michael. Machine Coaching. In: Proceedings of the IJCAI 2019 Workshop on Explainable Artificial Intelligence (XAI), pages 80–86, Macao SAR, P.R. China, August 2019.

3.10 Social biases: Identifying Stereotypes about Women and Immigrants – The Case of Misogyny

Paolo Rosso (Technical University of Valencia, ES)

Language has the power to reinforce stereotypes and project social biases onto others. At the core of the challenge is that it is rarely what is stated explicitly, but rather the implied meanings, that frame people's judgments about others. This is the case of stereotypes about women and immigrants, two social categories that are among the most preferred targets of hate speech and discrimination. In the first part of the talk, we address the problem of automatic detection and categorization of misogynous language in social media. Special emphasis is given to the issue of transparency during data collection and labelling, as well as at the time of the explanation of the categorization of misogynous language. Finally, we illustrate a taxonomy that we proposed to address the problem of stereotypes about immigrants.

References

1 Basile, V., Bosco, C., Fersini, E., Debora, N., Patti, V., Pardo, F. M. R., ... and Sanguinetti, M. (2019). Semeval-2019 task 5: Multilingual detection of hate speech against immigrants and women in twitter. In 13th International Workshop on Semantic Evaluation (pp. 54-63). Association for Computational Linguistics.
2 Fersini, E., Nozza, D., and Rosso, P. (2018). Overview of the evalita 2018 task on automatic misogyny identification (ami). EVALITA Evaluation of NLP and Speech Tools for Italian, 12, 59.
3 Sánchez-Junquera J., Chulvi B., Rosso P., Ponzetto S. (2021) How Do You Speak about Immigrants? Taxonomy and StereoImmigrants Dataset for Identifying Stereotypes about Immigrants. In: Applied Science, 11(8), 3610 https://doi.org/10.3390/app11083610

4 Working groups

4.1 Stakeholders and their Questions (day 1)

Judy Kay (The University of Sydney, AU), Tsvi Kuflik (Haifa University, IL), and Michael Rovatsos (University of Edinburgh, GB)

The online nature of the seminar necessitated a setup where sub-groups in compatible timezones convened for extended sessions twice each day, with those meeting in a larger session then further collaborating in small breakout groups on specific topics. The participants of each session were re-shuffled every day to maximise opportunities for all of them to interact with each other at least for part of the event. Further, to allow all participants to engage with the material presented, recordings of invited talks presented in each session were provided for the benefit of those not able to participate in those sessions due to their geographical location. The working groups across all sessions also collaborated asynchronously through shared online documents, documenting discussions and enabling cross-fertilisation between the work of individual groups.

This process was followed for the first three days of the seminar, with the final day being structured around working groups forming around topics for different sections of the joint "transprency by design" framework document authored jointly by all seminar attendees.

On the first day, working groups focused on a user-centric perspective to understand requirements for transparency. For this purpose, participants were split into small groups that aimed at identifying the questions users want to be answered when considering the use of algorithmic systems.

At a high level, these questions characterise the sorts of questions that a transparency by design approach will enable key stakeholders to answer. The group discussion was organised in six stages:

Stage 1 Select one key context that the group has some expertise and interest in, and identify the key classes of stakeholders.

Stage 2 For one such stakeholder group, we brainstorm to define an initial broad set of driving questions the people in that stakeholder group may want to be able to answer.

Stage 3 Repeat the above steps for a second stakeholder group.

Stage 4 Identify the questions that the group considers to be most important to support in the design process, taking account of additional stakeholders identified in Stage 1 – this was the narrowing phase of the brainstorm discussion.

Stage 5 Review the conclusions across the other groups of participants (reviewing the shared document used by all groups). Each group then used this to both refine the set of driving questions for that group's context and to identify the similarities and differences between them.

Stage 6 Finalize what to share with other groups.

The groups then returned from their breakout groups to a joint session, with a representative from each group presenting the key outcomes of their group's work. There were in total six groups and they selected three different domains, in the areas of education (3), job recruitment (2) and e-commerce (1).

For *job recruitment*, the identified key stakeholders are government regulators, certification agencies, users (job consultants, applicant (not) offered a position, lawyers contesting decisions, software developers, analysts, and hiring entities). The most important stakeholders interacting with a job recruitment system that were identified were regulators and applicants.

The key questions concerning the transparency of the system were “how”, “why” and “what”. There were two major points of view, applicant and job consultant, but the questions that were identified are relevant to both. For the “how”, these were: How is the candidate being scored? How is the user’s privacy protected? How much control do the users (employers/candidates) have to appeal/contest the decision? How well can the system handle discrimination bias? How can the candidate be a good fit for a specific job position? How does the system make a decision? For the “why”, they were: Why did the candidate not get the job position while someone else was successful? For the “what”, they were: What is the appropriate process followed for developing the software? What control do the end users have over the system (e.g. once we have identified that the system is doing something wrong, is the candidate able to change this)? What criteria were set for the system to operate? What concepts, values, terms have been defined in the system? What data was used to train the model? What process was used to make the decision to use that data? What are known historical good and bad practices? What went wrong in the past and why (no matter whether/what technology was used)? Have we investigated this history? What have we learned from these investigations?

For *education*, the key stakeholders are the students, teachers, professors, the parents, the university and school administrators, IT experts (technology developers and operators of the assessment system) at the institution, research funders and society at large (general public, journalists, civil society activists). The most important stakeholders identified in an educational system are the teachers and students. Again, the key questions were divided into “why”, “how” and “what”. The key questions concerning the transparency of the system are “why” questions, such as: Why does the system make an intervention (which can be justified based on the system’s objective), or why is this problem (specific topic/concept) recommended for the student to solve? In terms of “what”: What criteria does the system use to make a decision (e.g. Why did I get this grade? What are the marking criteria?)? For the “how”, the questions are about a variety of aspects regarding process and outcomes: How did the system select a peer group to compare the student’s performance to? How is a specific intervention generated? How would the grade differ compared to human-given grades? How can I improve my grade? How does the system produce the specific grade or decide the student’s level? How does an adaptive educational system determine/evaluate student knowledge for a domain topic or concept? How can we ensure consistency of assessment? How do I know that the accuracy of the assessment is correct? For the “what” the focus was on: What information is being shared with third parties (external vendors)?

In the *e-commerce* domain, the key stakeholders are buyers, sellers, creators – i.e. music track, performer vs. producer, advertisers, developers of a system, shipping/supply chain, recipients, regulatory (depending on product) operator. Here the main stakeholders are buyers and sellers and as they greatly differ, as do the questions. The key questions regarding the transparency of the system that will concern buyers are why they get (or do not get) specific recommendations? How to interface with the system, and how the system deals with any ethical and privacy concerns? How reliable/trustable the information and reviews of a product are that are presented? How people build their mental models of the space of products and services? The driving questions of the sellers concern the transparency of why their product is or is not recommended to specific buyers, how the pricing works, how the advertisement of products is performed, what responsibility issues there are; and what control a buyer or seller interacting with the system has. In general, the most important question is why someone gets a specific recommendation, what control buyers/sellers have while interacting with the system, and how trust in the system has been calibrated.

Through discussions in the working groups on the first day, a comprehensive account of requirements for transparency was elaborated for a number of different scenarios (including additional ones beyond those described above). These demonstrated very clearly that the breadth of possible concerns is potentially too great to condense into concise models, especially considering the specifics of different domains, and the societal, regulatory, legal, and ethical parameters that depend on the context of use of algorithmic systems in each sector. Nonetheless, it became clear that there are ways to systematically analyse who needs to know what about a system, when and how this information should be delivered to them, but also that these are issues that are not currently routinely addressed when algorithmic systems with potentially wide-ranging societal impacts are developed.

4.2 State of the Art and Emerging Priorities (day 2)

Judy Kay (The University of Sydney, AU), Tsvi Kuflik (Haifa University, IL), and Michael Rovatsos (University of Edinburgh, GB)

A major outcome of discussions on the first day of the seminar was the need to focus on collecting and analysing data in order to answer transparency-related questions. Framed within an exploration of the state of the art and research challenges around making progress in this direction, the second day followed a discussion format where working groups progressed through four stages: 1) Describing the state of the art on data collection in each of the domains discussed; 2) identifying gaps in terms of data collection; 3) considering requirements for an agile test-driven environment to bridge these gaps; and 4) identifying challenges that need to be addressed to enable such methodologies. The domains discussed in the six groups were similar to those of the previous day: e-commerce, job recruitment, and educational systems.

In the e-commerce domain, the main questions from the previous day were: Why did a customer get this specific recommendation, what control does the customer have over who sees the seller's product, and why are features to allow such customer-side control configured the way they are. Transparency in recommender systems is usually handled via providing explanations for recommendations in current systems, but many publications are also emerging that address fairness in recommendations. Looking at Amazon as a prime example, the user model can be built given based on short-term profile elements: current session interaction (items clicked, time spent on reading reviews/exploring a specific item, sequence of items), which is often combined with long-term data (the previous history of the user including user and item contexts, the behaviour of similar users – collaborative filtering – and derived attributes). Some of this data involves private self-reported information and some involves public data. Explicit data can also be collected from user ratings or surveys. Such data collection leads to additional questions regarding data transparency, e.g. how the specific user model is built and how it is used later on for recommendations, what training data is used, from what time period, whether different algorithms are used in different countries by the same company due to different legislation in those countries, etc. Some data pre-processing methods that can be used on the collected data such as adversarial inputs, detection of skew in available data for popular vs. less-popular items, data augmentation techniques and pre-trained embeddings for text or image analysis. The gaps from the existing literature regarding the data pre-processing techniques are that transparency issues of the data still exist especially when using pre-trained embeddings for text or image processing as well as a lack of transparency from technology providers regarding underspecification of how their systems operate, including how they may have been "sanity checked" and "stress tested".

Taking all these into consideration, and using an agile test-driven approach, one of the two groups suggested an approach where, at each software development cycle (requirements, design, implementation and testing) various stakeholders and technologies, appropriate for that specific stage, should be used. For example, for the specifications/requirements analysis phase the main stakeholders are end users, regulators and owners, and technologies/methods used are standard software engineering techniques. However, for a future transparency-driven development, where users would like to get both an explanation about the model itself and its outcomes, and to be able to see what features were used to build her actual user model, the following approaches should be considered: In the requirements stage, transparency could be defined as a non-functional requirement that might later become a functional requirement (for example, how to present the user profile/model to the user with features that were used for building it, e.g. via user preferences, with a checklist of relevant features the user can opt out of). This should be in addition to providing model explanations, as well as outcome (recommendation) explanations and global explanations for the overall software quality assessment.

In the testing stage, the generation of test cases to validate transparency requirements will be important (e.g. in regression testing), where the testing approach itself should also be transparent. In terms of stakeholders, it should be clear who is involved in each stage of the engineering process, what decisions they have to make, who is responsible for each decision, and whether these decisions need to be reported to regulators. Regarding the technologies used, those used during the certification processes should be independent of those used in development.

The group highlighted the importance of educating stakeholders in order to embed such an approach in systems development, but some open questions remain unresolved, including: What certifications are available in terms of trusting the system? Does the public want to be educated to understand such certifications? From an ethical/legal perspective, users need to know their rights. Were the engineers trained in transparent systems design, and how should this issue be addressed (e.g. should such training be legally mandated)?

Another group performed a literature survey on state-of-the-art research related to transparency in data collection in the area of recruitment. They identified papers and systems that deal with automatic recruitment and its implications, challenges, and solutions. Most of the papers referred to fairness in the recruitment systems, which in some sense overlaps with transparency. However, it is hard to identify applications implemented in the real world. In addition to this, some solutions proposed in the literature are not closely embedded in the context of hiring and recruitment, and are often more about generic machine learning based systems. For systems that deal with people and decisions taken about them, the questions seem to be similar. Moreover, there are generic solutions that exist and maybe we need to look into them and one should consider whether there is a need for adaptation towards objectives we might aim for in the recruitment domain.

In the context of educational systems, the questions for transparency concern mostly what kind of data the system has about students, how it ranks them, why a student obtained a particular recommendation, and how accurate this recommendation is (which is typically also asked by the educational provider). The group found many publications that include use cases such as machine learning systems that predict learning outcomes and forecast student grades. Another aspect discussed were privacy concerns around student data, as evident from two case studies (one on how Coursera handles GDPR issues, and another one using Named Entity Recognition to anonymize student data). In these case studies, the data used as input to algorithmic components includes demographic data, student grades, their responses in assignments, evaluations of teaching performance, represented either as text or in tabular data. Techniques that are usually used to pre-process the data included data cleaning techniques for the training data that would be used in a classification or clustering algorithm, web mining techniques, the use of association rules to mine learner profiles and diagnose learners' common misconceptions, sequential pattern mining to extract and present patterns that characterize the behavior of successful groups, and web mining techniques that group documents according to their topics and similarities and provide summaries. As for transparency, a case study that reports on using the LIME explainability method for student monitoring techniques used in a recommender system, and the QII approach for providing transparency reports in learning systems stood out as more principled studies in the state of the art.

Regarding the allocation of people to positions, job applications discussion focused on problems of both the regulators and the applicants, and, in particular, on helping the applicants to understand why they did not get a job. This is not only about looking for discrimination, but also about coming to a greater sense of self knowledge and job knowledge and and bringing all these pieces together. Since both regulators and naive users are involved in this process, an approach similar to digital forensics seems appropriate, not only for the regulators to ensure fair treatment, but also for the applicants who should be able to inspect differences between themselves and other, successful, applicants. Another important concern is personal data retention by platform providers, in particular how long personal data can be kept by the company and used. This is important not just because of GDPR, but also because of labor markets and social policy. Apart from understanding how an applicant is different from others, it is also essential to be able to identify whether there *are* appropriate jobs they can apply for with high chances of success, but also to help them understand how they can increase their chances. In other words, transparency should also be used to empower users to make their own appropriate choices rather than just explaining current systems behaviour.

Finally, in the college admissions context, further dimensions where discussed included explaining whether diversity targets in the admissions process are met and outcomes improving, understanding which students need/deserve more detailed explanations, and the wider issue of understanding which stakeholders require what different types of system output, including predictions and explanations. Concerning the data collection process, the system might use self-reported demographic and academic performance data that candidates enter in their applications, but input data might also include information from social media and the candidate's other online activities. Data pre-processing includes data cleaning and "munging", and also the use of feature engineering methods. In terms of data management, in state-of-the-art approaches, data provenance techniques are usually employed, and privacy policies may apply differently to different datasets and data sources. Open questions in this domain include: What are the important user characteristics and how are those measured, especially if they go beyond academic performance? Why is a particular piece of data being used and what is the evidence for it being relevant to the successful completion of the degree? The remaining gaps from the state-of-the-art concern the acquisition of all the necessary data to be in an accessible and usable format, and methods to assess the adequacy and validity of the system.

To summarise, from all the contexts chosen by the groups, there is a wide range of state-of-the-art publications on transparency in the respective domains. Concerning the data collection process, in almost all recommender systems, both implicit and explicit information are collected from users. Transparency is achieved using explanations for provided recommendations. In educational systems, the data collection process depends on the objective of the system but usually it concerns self-reported data (e.g. demographic data, responses to assessment questions) or information given by other users (e.g. student grades). Data pre-processing techniques might include the general data cleaning techniques or more specific techniques such as testing against adversarial inputs, web mining and data augmentation techniques. Also, various stakeholders of the system should be considered and explanations provided to them should differ to meet the particular needs of each stakeholder. In terms of implications for software engineering practice, the case studies demonstrate that transparency should be considered at the requirements phase and also throughout the whole software development lifecycle, which is not the norm in current practice. In particular, the shift to data-driven system mandates more rigorous planning, monitoring, and documentation of data provenance, collection, and analysis methods deployed.

4.3 Thematic Research Challenges (day 3)

Judy Kay (The University of Sydney, AU), Tsvi Kuflik (Haifa University, IL), and Michael Rovatsos (University of Edinburgh, GB)

Following on from working groups convened on day two, the third and final day of structured group work focused on mapping the existing research landscape to further identify gaps and seek to articulate an agenda for future research themes; with the aim of pulling results from all three days together to feed into a jointly authored programmatic paper that would provide the foundation for a "Transparency by Design" methodology in the subsequent final seminar day.

As before, working groups were asked to choose a concrete use case context to focus on, in order to identify those questions that are most pertinent to the chosen context, and explore state-of-the-art design and implementation methods which then lead to the identification of existing gaps. A specific focus for this day were challenges around building interfaces that can support users in scrutinising AI algorithms and pinpointing the challenges surrounding the creation of such interfaces. The discussions undertaken by the six groups covered some more general concerns on recommender systems as well as specific issues from one of the following domains: applicant ranking in an educational and employment context, e-commerce systems, care robots, and debugging tools for data-driven/AI systems, which are a cross-cutting concern for many domains.

Regarding systems that rank applicants and items, the following seem to be lacking: (a) an ability to trace the reasoning of deep learning systems, (b) the ability to explain to end users what features of their profiles were used for ranking, (c) the ability to opt in and out of the usage of certain features by the systems, (d) suggestions on what to do for lower-ranked users, and, as a more challenging gap, (e) the development of intrinsically transparent ranking methods, where transparency is not just provided *post hoc*. In the context of interfaces and explanations offered for rankings, common everyday life applications in social media

(Facebook), e-commerce (Amazon), video streaming (Netflix) and recruitment (LinkedIn) applications were discussed. In the example of video streaming platforms, services only show similar movies and percentages without explaining them. In terms of gaps, one major issue seems that it is impossible to find examples of interfaces for stakeholders other than the end users. Across many platforms, explanations seem to be very short and basic, and are likely to be insufficient for the information needs of at least some users, let alone other stakeholders such as analysts, regulators, and policy makers; where more detailed explanations also offer the opportunity to be open to end users as a side-benefit.

In the area of *e-commerce*, two main questions were identified: *Why am I getting a particular recommendation (consumer-side)?* and *Who sees my product (seller-side)?*. Explanation types vary across different types of data and the same explanation can be offered in different ways depending on the target audience. In this area, the core function of the system is to produce recommendations, so one way to look at the question is what type of explanations to use for this purpose. The areas of identified gaps included (1) systematic approaches and methodologies for designing transparency-compliant systems (e.g. training courses, regulations and other resources from relevant authorities, personalised explanations), (2) privacy and stereotyping related issues (e.g. understanding the trade-off between privacy and personalised explanations), (3) auditing and certification (e.g. regular and ongoing auditing). When it comes to user interfaces there is a lack of multimodal explanations and discoverability of existing explanations. From an ethical point of view, it is important to ask how not to abuse explanations and the interface to manipulate people (or for the platform to be manipulated by its users).

In the field of *care robots*, the main questions identified were: What is the aim of building such technology, who is meant to be the beneficiary, who are the explanations for, how are decisions taken by the robot (which can change the relationship between the carer and the patient) and what level of transparency is needed during the development, data collection, certification phase and deployment of such robots? Multiple potential areas of transparency were discussed, such as the transparency of the recommendations provided by the robot (i.e. its reasoning behind concrete decisions), transparency of how the system exactly operates and transparency for regulatory purposes. It was suggested that different levels of transparency are needed for different areas but participants noted that in certain cases transparency can lead to privacy concerns. For instance, capturing the original design and development conversations can be useful in order to demonstrate due diligence but also can endanger privacy and cybersecurity of the users of the robots. The capturing of abstracted but comprehensible levels of development and operations information is necessary for the maintenance, inspection, installation and explaining of the system. Since care robot systems give advice on high-stake matters, it is very important that the advice can be traced back to the development process as well as the medical literature and research underpinning their design.

The discussions in the working groups on this day revealed that the landscape of methods proposed by the academic literature on various sub-topics related to transparency is indeed very rich, and ranges from new advances to improve the explainability properties of AI algorithms and fairness-aware recommender system design all the way to work on advanced user interfaces that enhance the understandability and configurability of algorithmic systems. However, the current literature not only presents itself as fragmented in terms of providing complete accounts of end-to-end transparency-driven approaches to systems development, there is also a (much more serious) lack of connecting the specific research advances to actual systems in the real world, where very little is known about how they are actually implemented. Given this, it is unsurprising that some of the existing work seems hard to apply to real-world examples, with authors making often unrealistic assumptions that can only be fulfilled in experimental systems developed purely for research purposes. Overall, it seems that, while some technical challenges remain, the methodological building blocks for developing a holistic transparency by design approach are largely in place. What is missing is an ability to map them onto a context-specific methodology that could be embedded in the development of impactful (usually commercial) real-world systems. In this regard, the most important roadblock as the moment is an inability to evaluate these systems to propose concrete transparency improvements. There are significant barriers to overcome here in terms of cross-sector collaboration and education, and our overall longer-term objective of proposing a methodological framework might help accelerate the dialogue between researchers, developers, regulators, and users in this respect.

4.4 Toward a Transparency by Design Framework (day 4)

Judy Kay (The University of Sydney, AU), Tsvi Kuflik (Haifa University, IL), and Michael Rovatsos (University of Edinburgh, GB)

During the fourth and final day, the results of the first three days were discussed and summarised into a joint document that is intended to form a basis for a joint paper. The structure of the document followed the results of the discussion of the topics and the order of the discussion in the first three days, organised into the following chapters:

1. Why transparency?
2. Challenges posed by current algorithmic systems
3. Transparency-enhancing technologies
4. Transparency by design – principles
5. Transparency by design – methodology
6. The Road Ahead

Throughout the day, and also after the conclusion of the seminar, self-selecting sub-groups of seminar participants collaboratively developed material for each individual section, working synchronously and asynchronously. The work of these groups was guided by an overarching structure and guidelines for for the content of the planned paper, following a set of key questions that guided the discussion, which we list below together with key results of the discussions:

- Key questions to answer about “Why Transparency”:
  1. What is transparency and why is it important?
  2. How does it relate to other concerns around AI?
  3. Why is it timely to develop new transparency capabilities?

The high-level results of this discussion included that there are many public policy documents that call for transparency, but that different experts take very different perspectives and talk about different things when it comes to transparency. Nonetheless, there is some universal agreement on the importance of transparency for trust and its calibration, in order to have accurate mental models of systems for users, to support accountability and auditing, and to enhance fairness and human control. It was noted that transparency should not be taken as a panacea, or an end in itself. Rather, transparency can contribute to ensuring human accountability, whether this comes from the use of good design and development practices, addressing –or even readdressing– problem issues, or to underpin appropriate levels of legal liability.

- Key questions to answer about "Challenges posed by current algorithmic systems":
  1. What has changed in current systems that creates challenges for transparency?
  2. How hard will it be to enable pervasive transparency given these challenges?
  3. Why does this matter to users, organisations, societies?

The views discussed here highlighted that the key changes that create new challenges are due to a growth in the ubiquity of AI-based systems, the presence of particular impacts of these system on vulnerable groups and the growing cost these incur, but also the increasing availability of tools for creating these systems, which accelerate these impacts. This can be seen as a consequence of the "democratization" of AI, which makes it so readily available that it is incorporated in many systems. Furthermore, data acquisition has deeply changed over time, from a time when the AI system did not know anything that was not inserted in it by a human directly to modern-day machine learning using global-scale Internet data, which makes it harder to audit for and filter content, with the real risk of poor data and associated inferences. A number of increasingly prevalent practices can give rise to further transparency risks: The use and reuse, including sometime on a large scale, of datasets that lack proper documentation; the development of "general" pre-trained models as "cognitive services", which developers can use in a plug-and-play manner; and the emergence of application-agnostic components and tools released in the public domain without knowing or being able to control their downstream usage. On a societal level, there is a growing demand for accountability at the level that would be expected of people making the decisions that are entrusted to systems, and this is reflected in emerging legal and regulatory requirements, together with a growing recognition of needs and challenges of education about AI, including AI literacy of the general public and curricula for schools. Participants identified many factors that make it challenging for developers to support transparency, including: The prevailing attitudes of those responsible for major players; the technical difficulty of even defining what we mean by transparency, much less enabling it; challenges related to designing for multiple stakeholders; the potential for adversarial interaction to game systems that might arise from increased transparency; and a lack of understanding of the complex dynamics of human-AI systems. This presents us with a formidable challenge, because of the many costs of the current lack of transparency, which may entrench disadvantage and inequality through a broad range of unanticipated societal impacts. It was noted that it is extremely hard to find general solutions to transparency problems. Domain-specific analysis is key to be able to do solid assessments of potential solutions specific to the use case. This is particularly important as, more often than not, the multiple stakeholders have different conflicting objectives.

- Key questions to answer about "Transparency-enhancing technologies":
  1. What approaches to enhancing transparency have been proposed?
  2. What benefits did they deliver, and what were their limitations?
  3. What are the transparency issues they did not address?

In the discussion of these questions, participants focused on several issues. The first one is considering all stages of the design and development process of AI-based systems, including data preparation, the learning algorithm/model/process used, model and outcome explanations, and user interfaces through which these will provided. Along a second dimension, each of this issues needs to be considered for all relevant stakeholder groups, from the model and software developers to project managers and end users. This needs to be done in a fine-grained way, as, for example, different types of end users may have very different requirements. The participants discussed a range of methods available (and missing) from the literature, highlighting the importance of distinguishing between different types of data (e.g. text, images, video, etc), its source (e.g. offline data collection vs. crowdsourcing), the annotation processes used, traceability and provenance methods used, fairness metrics applied. They discussed differences between types of learning algorithms used, e.g. supervised, semi-supervised, unsupervised, and reinforcement learning methods, and to what extent different explainability methods can be applied to them. For the latter, there already exist useful taxonomies of explainability methods (including rule-based white-box models, surrogate models of black-box models, and counterfactual explanations), and there is evidence that users favour "why" and "how" explanations along general lines, but that one-size-fits-all explanations are not always effective or desirable. Finally, in terms of interfaces, there is a body of work on these especially in the area of recommender systems, which include visualisation of latent influencing data and variables, the explication of intermediate results, but also new interfaces that allow improved user control and scrutability.

- Key questions to answer about "Transparency by design – principles":
  1. What are the high-level guiding principles we want to embed in Transparency by Design?
  2. What is the justification behind them? (Why do we think they are important?)
  3. How realistic is it that they could be achieved? (Is this aspirational or practical?)

Under this theme, the participants broke down the overall question into issues relating to defining transparency by design in terms of its scope, purpose, and method. The first question here is to consider where transparency should be enabled, focusing on the design process, which should explicitly state assumptions and context, trace the provenance of data used and how it flows through a system, make the construction process of models and the use of their outcomes explicit, and document testing and validation processes and outcomes. The second is clarifying the purpose of transparency, which should be to empower users (enabling contestability, data subject rights enforement, configurability of systems, and their risk management). Transparency by design approaches will need to effectively demonstrate they address the requirements of different stakeholders in this regard, including customers, regulators, developers, system integrators. In term of the methods underpinning transparency, a framework will have to specify how they are evaluated and validated (e.g. through benchmark tasks and metrics), establish principles for ensuring that the relevant bodies and individuals will and can understand the information provided, and provide a normative framework for understanding when explanations are appropriate, who is responsible for providing it, and what the appropriate ways of communicating the information are.

- Key questions to answer about "Transparency by design – methodology":
  1. What changes need to be made in each stage to satisfy transparency principles?
  2. Who are the people that need to take responsibility for applying these techniques?
  3. Can we describe these techniques through a relatively simple process model?

The high level results of the discussion in this area reiterated the importance of taking a holistic approach, where transparency would apply to input data, learning process, decision making, and output. A solid methodology would include a full account of decisions made by the system (or by humans using it) through a range of explanation methods using transparency-enhancing user interfaces. Participants proposed a preliminary "five star" framework of transparency, building on the following pillars: (1) explanation of decisions; (2) stakeholder-appropriate provision of these; (3) user interfaces that support transparency and exploration of the system; (4) open sharing of validation process and results; and (5) comprehensive description of the input data.

- Key questions to answer about "The road ahead":
  1. What are key problems the research community needs to focus on?
  2. What kinds of experts (academic/industry/government) are needed to do this work?
  3. What other enabling steps need to happen, e.g. in industry or government?
  4. What are the limits of TBD, what problems will it not solve?

For this final them, the groups discussed and contributed insights on suitable ways to further pursue a Transparency by Design agenda. One important aspect that was raised here was the importance of advancing general principles for corporate, legal, and product transparency in AI. This will require enhancing the wider understanding of existing regulatory obligations, outreach and education to legislators, enforcers, civil society watchdogs, judges, developers and tech corporations. Funding and political will to enforce societal expectation will be important, together with substantial efforts to counter misinformation emanating from AI-based systems, which is recognised as one of the major problem areas. It is important to acknowledge existing initiatives, including proposals for regulatory frameworks and standards. Undoubtedly, a lot of further development work and research will be necessary, which might even lead to new types of professional roles for people supporting AI transparency with responsibilities for oversight, monitoring, and enforcement, embedded in different types of organisations. Before effective transparency capabilities are developed across the industry, we will likely see much more work on standards, exchange of good practice, the resolution of tensions between transparency and intellectual property rights, and, of course, technical development of new transparency-enhanding technologies and interfaces.

Participants

- Elisabeth André
Universität Augsburg, DE
- Bettina Berendt
TU Berlin, DE
- Nehal Bhuta
University of Edinburgh, GB
- Maria Bielikova
KInIT – Bratislava, SK
- Veronika Bogina
University of Haifa, IL
- Joanna J. Bryson
Hertie School of Governance – Berlin, DE
- Robin Burke
University of Colorado – Boulder, US
- Aylin Caliskan
George Washington University – Washington, DC, US
- Carlos Castillo
UPF – Barcelona, ES
- Ewa Cepil
UPJPII – Krakow, PL
- Cristina Conati
University of British Columbia – Vancouver, CA
- Aviva de Groot
Tilburg University, NL
- Gianluca Demartini
The University of Queensland – Brisbane, AU
- Virginia Dignum
University of Umeå, SE
- Fausto Giunchiglia
University of Trento, IT
- Riccardo Guidotti
University of Pisa, IT
- Meeri Haataja
Saidot Ltd – Espoo, FI
- Judy Kay
The University of Sydney, AU
- Styliani Kleanthous
Open University of Cyprus – Nicosia, CY
- Ansgar Koene
EY Global – London, GB
- Joshua A. Kroll
Naval Postgraduate School – Monterey, US
- Antonio Krüger
DFKI – Saarbrücken, DE
- Tsvi Kuflik
Haifa University, IL
- Bob Kummerfeld
The University of Sydney, AU
- Loizos Michael
Open University of Cyprus – Nicosia, CY
- Antonija Mitrovic
University of Canterbury – Christchurch, NZ
- Kalia Orphanou
Open University of Cyprus – Nicosia, CY
- Jahna Otterbacher
Open University of Cyprus – Nicosia, CY
- Anna Perini
CIT- FBK – Povo, IT
- Lena Podoletz
University of Edinburgh, GB
- Iris Reinhartz-Berger
University of Haifa, IL
- Paolo Rosso
Technical University of Valencia, ES
- Michael Rovatsos
University of Edinburgh, GB
- Avital Shulner-Tal
University of Haifa, IL
- Alison Smith-Renner
University of Maryland – College Park, US
- Andreas Theodorou
University of Umeå, SE
- Vincent Wade
Trinity College Dublin, IE
- Emine Yilmaz
University College London, GB

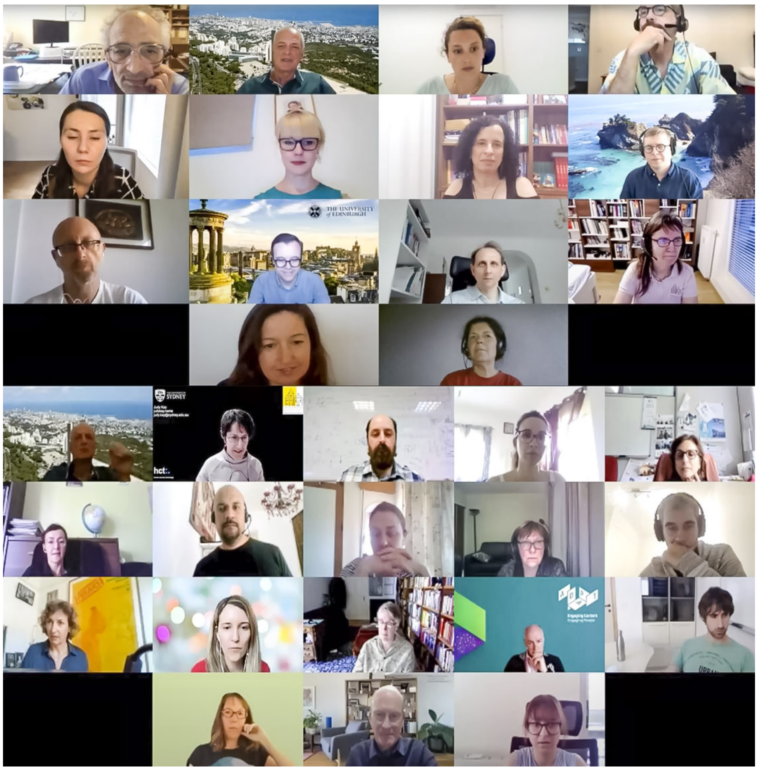

Report from Dagstuhl Seminar 21232

Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments

Edited by

Stephen Brewster[1], Andrew Kun[2], Andreas Riener[3], and Orit Shaer[4]

1 **University of Glasgow, GB,** `stephen.brewster@glasgow.ac.uk`

2 **University of New Hampshire – Durham, US,** `andrew.kun@unh.edu`

3 **TH Ingolstadt, DE,** `andreas.riener@thi.de`

4 **Wellesley College, US,** `oshaer@wellesley.edu`

Abstract

We are living in a world where ubiquitous computing devices are becoming parts of the fabric of our lives. At work and at school, devices such as calculators, tablet computers, mobile phones, and different electronic measurement devices, support our work and learning. Building on all of these technological advancements will be novel human-computer interaction techniques that will allow us to use the devices for work and play in a broad set of circumstances, from riding in automated vehicles, to exploring museums, to walking on the street, to playing with our kids on the beach. The central underlying question Dagstuhl seminar 21232 wanted to address is, "how will we interact with the ubiquitous devices of our near (and not-so-near) future?" To date, there are a number of interaction techniques that show significant promise, including speech, augmented reality, tangible objects, gesture, multitouch screens, as well as simple keyboards and non-touch displays. But, before we address technologies to use, we must first identify the economic and broad societal driving forces that will create the need for interaction with our ubiquitous computing devices. From the economic point of view worker well-being is one such driving force; another one is the need to improve the productivity of workers and firms; yet another is the need to provide access to continuous education to a changing workforce. From the broad perspective of our society, it is important for us to understand how ubiquitous technologies can support living a meaningful and fulfilling life, from childhood to adulthood.

In the following, we report the program, activities, and the outcomes of Dagstuhl Seminar 21232 "Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments".

Seminar June 6–11, 2021 – http://www.dagstuhl.de/21232

2012 ACM Subject Classification Human-centered computing → HCI theory, concepts and models; Human-centered computing → Interaction paradigms

Keywords and phrases (Productive) Work, Ergonomics, Human-computer interaction, Wellbeing

Digital Object Identifier 10.4230/DagRep.11.5.23

Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments, *Dagstuhl Reports*, Vol. 11, Issue 05, pp. 23–53

Editors: Stephen Brewster, Andrew Kun, Andreas Riener, and Orit Shaer

Dagstuhl Reports

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Executive Summary

Andreas Riener (TH Ingolstadt, DE)
Stephen Brewster (University of Glasgow, GB)
Andrew Kun (University of New Hampshire – Durham, US)
Orit Shaer (Wellesley College, US)

Agenda in a nutshell

The seminar was conducted online during the week of June 6-10, 2021. A particular difficulty in planning the agenda (see Fig. 1) arose due to the different time zones of the individual participants. It was especially important to us to offer at least some of the program items together to all participants (Opening and Group Work on Day 1, Summit and Closing on the last day). On the other days, we planned different activities in smaller (2-3 people) and larger (up to half of the participants) groups to better accommodate the participants based on their time zones.

Dagstuhl Seminar 21232 "Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments"

					all	all	individual/small groups	individual/small groups	all
Europe (CEST, Germany)	US (EDT, e.g. Boston MA)	US (West, Seattle SA)	China/Hong Kong	Tokyo, Japan / Australia	Monday, June 7	Tuesday, June 8	Wednesday, June 9	Thursday, June 10	Friday, June 10
12:00:00 am	6:00:00 pm	3:00:00 pm	6:00:00 am	7:00:00 am					
1:00:00 am	7:00:00 pm	4:00:00 pm	7:00:00 am	8:00:00 am				Video discussions (pairs)	
2:00:00 am	8:00:00 pm	5:00:00 pm	8:00:00 am	9:00:00 am				Compilation of playlists (pairs)	
3:00:00 am	9:00:00 pm	6:00:00 pm	9:00:00 am	10:00:00 am					
4:00:00 am	10:00:00 pm	7:00:00 pm	10:00:00 am	11:00:00 am					
5:00:00 am	11:00:00 pm	8:00:00 pm	11:00:00 am	12:00:00 pm					
6:00:00 am	12:00:00 am	9:00:00 pm	12:00:00 pm	1:00:00 pm					
7:00:00 am	1:00:00 am	10:00:00 pm	1:00:00 pm	2:00:00 pm					
8:00:00 am	2:00:00 am	11:00:00 pm	2:00:00 pm	3:00:00 pm					
9:00:00 am	3:00:00 am	12:00:00 am	3:00:00 pm	4:00:00 pm		Workshop1 G1 (19)			
10:00:00 am	4:00:00 am	1:00:00 am	4:00:00 pm	5:00:00 pm		Workshop1 G1	Movie "Coded bias" +discussion (Miro)		
11:00:00 am	5:00:00 am	2:00:00 am	5:00:00 pm	6:00:00 pm		Workshop1 G1	Movie "Coded bias" +discussion (Miro)		
12:00:00 pm	**6:00:00 am**	**3:00:00 am**	**6:00:00 pm**	**7:00:00 pm**					
1:00:00 pm	7:00:00 am	4:00:00 am	7:00:00 pm	8:00:00 pm					
2:00:00 pm	8:00:00 am	5:00:00 am	8:00:00 pm	9:00:00 pm					
3:00:00 pm	9:00:00 am	6:00:00 am	9:00:00 pm	10:00:00 pm					
4:00:00 pm	10:00:00 am	7:00:00 am	10:00:00 pm	11:00:00 pm					
5:00:00 pm	11:00:00 am	8:00:00 am	11:00:00 pm	12:00:00 am	Opening & Introduction	Workshop2 G2 (US, 11)			Summit
6:00:00 pm	12:00:00 pm	9:00:00 am	12:00:00 am	1:00:00 am	Opening & Introduction	Workshop2 G2 (US)	Movie "Coded bias" +discussion (Miro)		Summit
7:00:00 pm	1:00:00 pm	10:00:00 am	1:00:00 am	2:00:00 am	Opening & Introduction	Workshop2 G2 (US)	Movie "Coded bias" +discussion (Miro)		Summit
8:00:00 pm	2:00:00 pm	11:00:00 am	2:00:00 am	3:00:00 am					
9:00:00 pm	3:00:00 pm	12:00:00 pm	3:00:00 am	4:00:00 am					
10:00:00 pm	4:00:00 pm	1:00:00 pm	4:00:00 am	5:00:00 am					
11:00:00 pm	5:00:00 pm	2:00:00 pm	5:00:00 am	6:00:00 am					

Figure 1 Compact overview of the agenda for the week including different geographical zones (for better planning with participants from all-over the world).

- Monday, June 6: The seminar was opened and its main goals introduced by the seminar co-organizers Stephen Brewster, Andrew Kun, Andreas Riener and Orit Shaer. The presented slides can be accessed here: `https://docs.google.com/presentation/d/15NtQy96wAS_dMHpdqTO-2TfZhbSGssWAn96RxHJizVA/edit#slide=id.gdd9402fdbb_0_8`. After a social "warm-up" activity, Pecha Kucha presentations of all participants followed. During the presentations, all participants were instructed to collect questions, ideas, thoughts, etc. on a Miro-board; The items were clustered by the organizers (in a short coffee break) and after that, a voting of topics to be picked-up/focusing on in the next days of the seminar (see Fig. 2) followed. This activity ended day 1.
- Tuesday, June 7: The second day of the seminar was dedicated to the "Work(shop) for the Future of Work and Mobility in Automated Vehicles". In this workshop, participants (see Fig. 3) worked together on user needs and how to fulfill them during shared or private automated mobility. The workshop was conducted twice – each with half of the

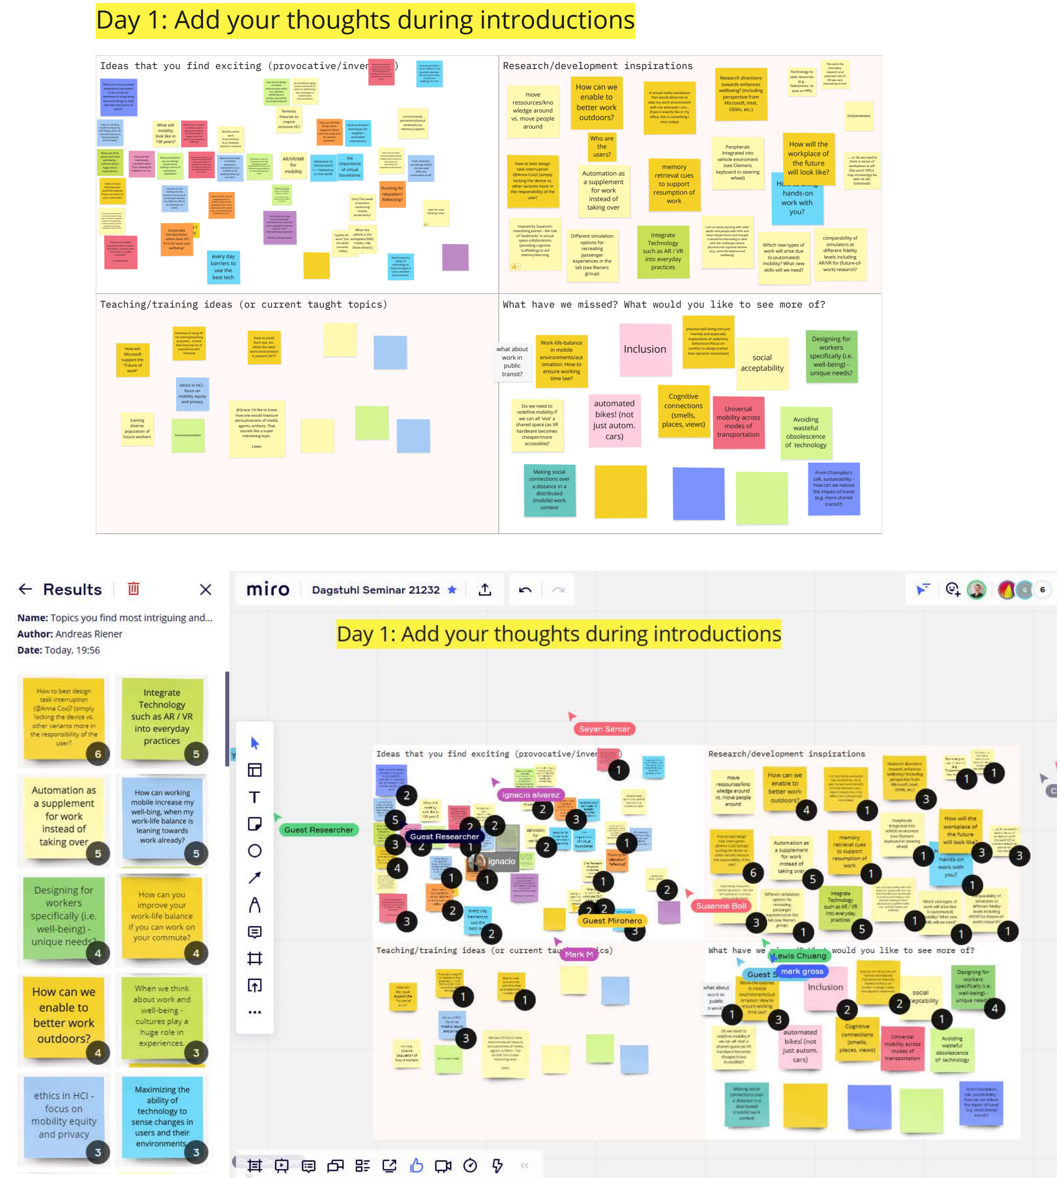

Figure 2 Group activities on day 1: Collecting of ideas, thoughts, questions from the individual presentations; Majority voting after clustering of collected items.

participants and lasted for about two hours including a short coffee break. In order to get all participants in the mood for the workshop and to allow them to reflect on the topic from their personal point of view, we invited everybody to complete a brief (10 min.) "pre-questionnaire" before the workshop (Link: `https://thimib.fra1.qualtrics.com/jfe/form/SV_03eUqLNatcDgzs2`). For details, see section 3.3. The results from both the questionnaire and the two workshops are currently analyzed and will be later submitted as conference paper or journal article (with recognition of the Dagstuhl seminar).

- Wednesday, June 8: On this day, in the Dagstuhl tradition to offer a social activity, we watched – again in two groups of each ca. 15 people – the documentary "Coded Bias" (`https://www.codedbias.com/`). While watching the video, participants were asked to record their thoughts (issues, concerns, suprises, technical problems/solutions,

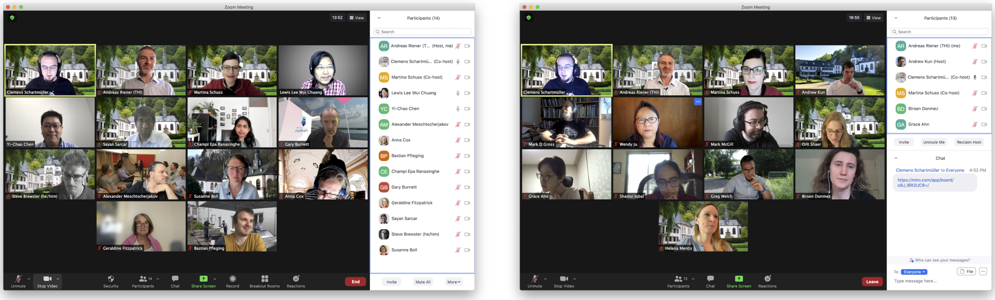

Figure 3 Introduction to the two workshops on day 2 including participants.

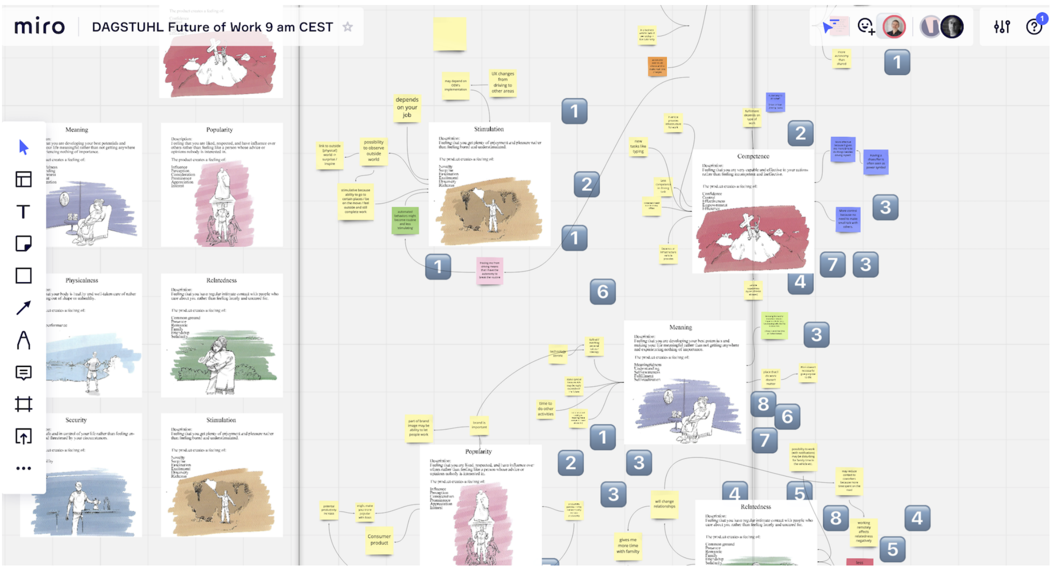

Figure 4 Intermediary results of the interactive workshop part on the Miro-board (group 2 workshop).

societal/policy related solutions) in a Miro-board, e.g., `https://miro.com/app/board/o9J_lCMqEEY=/` for group 1. After watching, we used 10 minutes for clustering the items followed by another 5 minutes for voting. The top voted items where than discussed in the large group and conclusions drawn for our work.

Coded Bias – group 1 results:

- 5 votes: "ensure the right to be forgotten" (removal/deletion of data)
- 4 votes: "AI algorithm uses historical information for the prediction – not everything has been seen before..."
- 3 votes: "Salery automatically based on office environment (stationary, in the car, on the go) -> lot of discussion
- 2 votes: "Transparency"
- 2 votes: "Use a diverse data set to train the AI"
- 2 votes: "Ways of opening the black box...?"

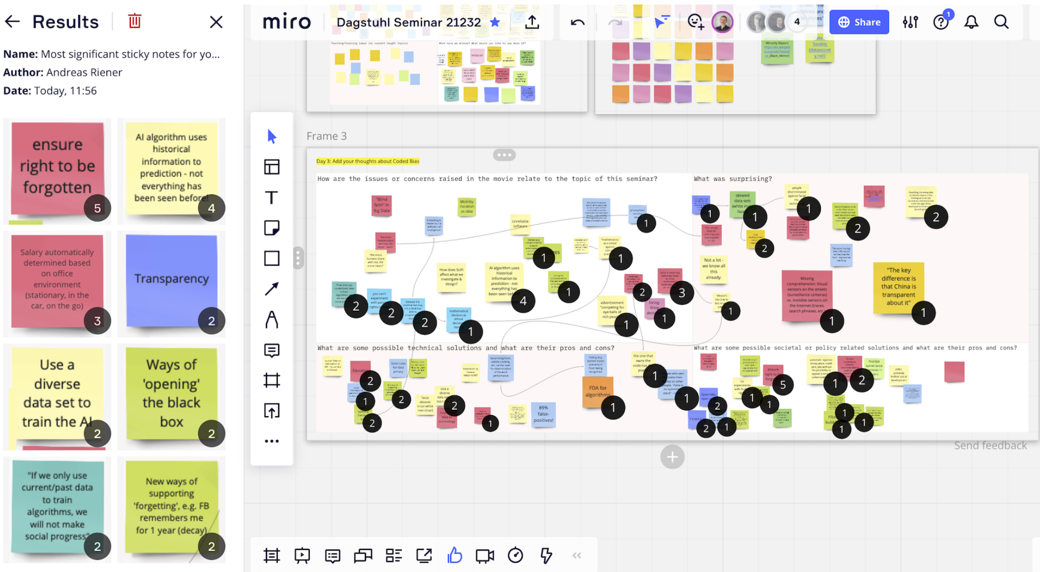

Figure 5 Post-its collected by the participants of group 1 and voting results.

Coded Bias – group 2 results:

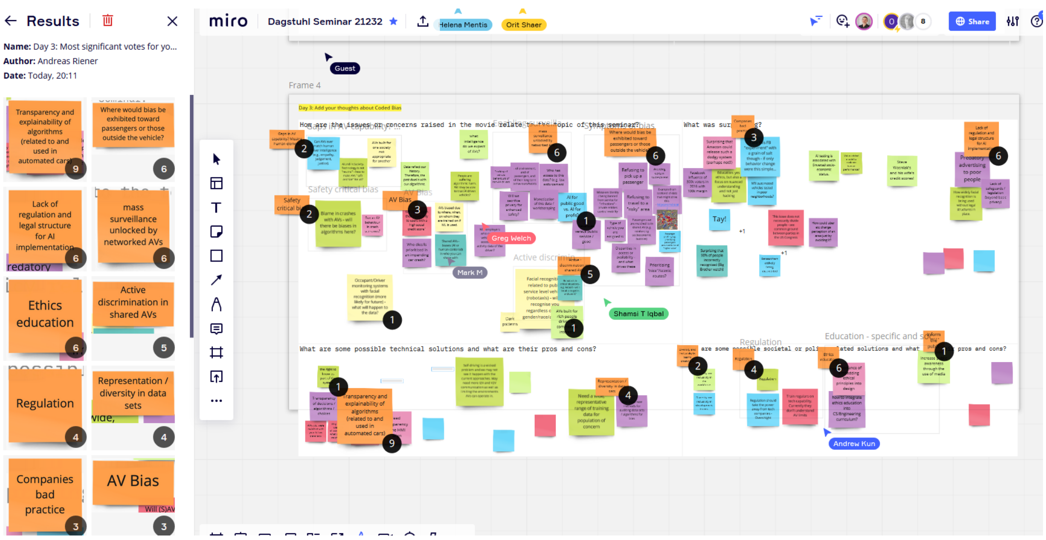

Figure 6 Post-its collected by the participants of group 2 and results of the voting on most relevant elements identified during watching "Coded Bias".

- 9 votes: "Transparency and explainability of algorithms (related to and used in automated cars)"
- 6 votes: "Where would bias be exhibited toward passengers or those outside the vehicle?"
- 6 votes: "Lack of regulation and legal structure for AI implementation"
- 6 votes: "mass surveillance unlocked by networked AVs"
- 6 votes: "Ethics education"

- Thursday, June 9: On the second last day of the seminar, all seminar participants met in small groups (2 to 3 people, see Fig. 7) to discuss one of the topics identified as most important (and to make a video of the discussion) or to jointly create a Youtube playlist of most-impactful videos in a dedicated topical area related to the seminar. The results were collected by the co-organizers of the seminar and distributed among the participants. Examples of bilateral interviews can be found in Sections 4.1 or 4.2, among others, and an example of a playlist is shown in Section 4.3.

Dagstuhl Seminar 21232 "Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments"

Participant 1	Affiliation	Participant 2	Affiliation
Chen, Yi-Chao	Shanghai Jiao Tong University	Fitzpatrick, Geraldine	TU Wien
Riener, Andreas	TH Ingolstadt	Burnett, Gary	University of Nottingham
Schartmüller, Clemens	TH Ingolstadt	Ranasinghe, Champika Manel Epa	University of Twente – Enschede
Shaer, Orit	Wellesley College	Lee, John D.	University of Wisconsin – Madison
Boffi, Laura	University of Ferrara	Brewster, Stephen	University of Glasgow
Ahn, Sun Joo	University of Georgia – Athens	Welch, Gregory F.	University of Central Florida – Orlando
Alvarez, Ignacio J.	Intel – Hillsboro	Gross, Mark D.	University of Colorado – Boulder
Boll, Susanne	Universität Oldenburg	McGill, Mark	University of Glasgow
Brumby, Duncan	University College London	Pfleging, Bastian	TU Eindhoven
Meschtscherjakov, Alexander	Universität Salzburg	Cox, Anna	University College London
Chuang, Lewis	IfADo – Dortmund	Lindley, Siân	Microsoft Research – Cambridge
Mentis, Helena M.	University of Maryland – Baltimore County	Ju, Wendy	Cornell Tech – New York
Kun, Andrew	University of New Hampshire – Durham		Sayan Sarkar, Tsukuba University
Donmez, Birsen	University of Toronto	Iqbal, Shamsi Tamara	Microsoft Research – Redmond

Figure 7 Couples who either had a curated conversation or created a Youtube playlist on Thursday bilaterally (<= 5 minutes each).

- Friday, June 10: The last day of the seminar has ended with a summit (Fig. 8). The first half of this activity was devoted to two panels with distinguished panelists. Panelists started the conversations with brief statements, which were then followed by moderated discussions with the group. For the second half of this activity all participants were sent into breakout rooms in Zoom and worked in smaller groups on a Miro-board (`https://miro.com/app/board/o9J_l_-usxU=`) on problems discussed during the panels. After the group work, all met again in the main Zoom room and each group presented the results of the group activity (Fig. 9).

Figure 8 The highlight of the seminar: A summit with contributions from seminar participants and keynote speeches from invited experts (including Neha Kumar, ACM SIGCHI President).

Figure 9 Overview of the results of the six groups in Miro.

2 Table of Contents

3 Overview of Talks

3.1 My Personal Research Outlook

Sun Joo Ahn (University of Georgia – Athens, US)

The future of work is anticipated to be a tight integration of workers and machines in both virtual and augmented versions of reality. Physical and geographical distance will become less relevant as workers learn to adapt to digital platforms that provide shared virtual spaces for social interaction and communication. Research on how these developments may transform the future of work has typically focused on worker efficiency but worker wellbeing needs to take into consideration the social relationships that support workers' resilience (e.g., family interactions [1]) and the context in which work takes place (e.g., working from home [2]). Users will begin to form relationships with artificial agents that drive and represent autonomous vehicles [3], and these relationships will impact how users interact with the vehicles. Future work in this area must consider individual, social, and environmental variables of the human-computer interactions that take place in the future of work, moving above and beyond the current focus on worker efficiency.

References

1 Rukangu, A., Mattingly, K., Franzluebbers, A., Tuttle, A., O'Neal, C., Robinson, D., Ahn, S. J., Johnsen, K. *Virtual family room: Bridging the physical distance with virtual reality*. Proceedings of Interservice/Industry Training, Simulation, and Education Conference, Florida, USA, 2020

2 Ahn, S. J., Levy, L., Eden, A. L., Won, A. S., MacIntyre, B., & Johnsen, K. J. *IEEEVR2020: Exploring the first steps toward standalone virtual conferences.* . Frontiers in Virtual Reality – Technologies for VR, 2021

3 Ahn, S. J., Johnsen, K., Ball, C.*Points-based reward systems in gamification impact children's physical activity strategies and psychological needs*. Health Education & Behavior, 46(3), 417-425, 2019

3.2 Cars with an Intent

Laura Boffi (University of Ferrara, IT)

Since its early phases, the development of automated vehicles has focused on functionality, safety and efficiency, overlooking the social implications that such technology would bring up in people's lives. "Cars with an Intent" is a design driven PhD research project which envisions how autonomous cars' behaviour and services can enable new car-to-human and human-to-human relationships. In particular, the research focuses on the "Co-Drive" concept, which is proposed as a new service for traveling and socializing by car between a driver of an automated vehicle and a remote passenger connected via virtual reality from home. I argue that the convergence of automated driving and telepresence technologies could provide a new social context for personal interactions to emerge, that are neither dependent on any earlier relationships nor based on age affinities. As remote passengers will likely be elderly people and drivers younger ones, the aim of the research is to understand how the "Co-Drive" concept could support intergenerational encounters and relationships and reduce the sense of loneliness in senior adults. Moreover, I aim to understand how the "Co-Drive" service could enable an ageism-free approach towards senior adults through the use of digital and robotic embodiment for remote passengers. Prospective remote passengers can select a location from where to start their Co-Drive trip from among the many stops around the world which have been featured in the Co-Drive Atlas and place their avatar there to book their trip. A Co-Drive stop is the physical location where a Co- Drive trip can start from, blending the real and virtual into an extended reality experience for the driver as well as for the remote passenger. A driver passing nearby the stop where an avatar has been placed would spot such an avatar as an AR visualization on the car windshield. S/he and can decide to pull over, start engaging remotely with the person embodied in the avatar and eventually board her/him in the car as a remote passenger for a shared trip together. They could both converse "live" during the trip, while the remote passenger could also enjoy the view as if sitting on the passenger seat. At the moment I am experimenting with different immersive system for the remote passenger: from a computer display, to a big wall projection to a VR headset. The project has followed a participatory design process, engaging with real participants in their context since the very beginning. I have been designing my own early-stage XR prototyping methodology iand, beyond crafting lo-fi artefacts, I have been establishing collaboration with external partners which could provide me with "enough technology" to pilot the remote trips, such as Ericsson R&D Italy which developed the car-pod.

References

1 Laura Boffi, Giuseppe Mincolelli, Simone Bertucci, Lorenzo Gammarota, Fabio Pes, and Marco Garofoli. 2021. Co-Drive: the experience of a shared car trip between a driver and a remote passenger. In ISMAR 2021proceedings, October 2021, Bari, Italy. [Forthcoming]

2 Laura Boffi. 2021. Designing for place-making in XR: the process of the Co-Drive stops and its atlas. In Media Architecture Biennale 20 (MAB20), June 28-July 2, 2021, Amsterdam and Utrecht, Netherlands. ACM, New York, NY, USA, 5 pages. https://doi.org/10.1145/3469410.3469434

3 Laura Boffi, Giuseppe Mincolelli, Simone Bertucci, Lorenzo Gammarota, Fabio Pes, and Marco Garofoli. 2021. Co-Drive: Experiencing Social Virtual Travel on a Car Trip. In CHI Conference on Human Factors in Computing Systems Extended Abstracts (CHI '21 Extended Abstracts), May 8–13, 2021, Yokohama, Japan. ACM, New York, NY, USA.

4 Laura Boffi, Philipp Wintersberger, Paola Cesaretti, Giuseppe Mincolelli, and Andreas Riener. 2019. The first co-drive experience prototype. In Proceedings of the 11th International Conference on Automotive User Interfaces and Interactive Vehicular Applications: Adjunct Proceedings (AutomotiveUI '19). ACM, New York, NY, USA, 254-259. DOI: https://doi.org/10.1145/3349263.3351318

3.3 Multimodal (and more physical) Interaction to Improve In-Car User Experience

Champika Ranasinghe (University of Twente – Enschede, NL)

Automated vehicles have the potential to provide increased mobility to a broader range of users such as elderly, children, people with physical limitations (e.g., visually impaired people, people with broken legs/ arms) or people with other types of limitations such as for example people who are nervous of driving [1] [2]. On the other hand, driving becomes a secondary task and the drivers (and the passengers) can engage in various other tasks such as work, leisure or even sleep. This requires interacting with the vehicle at various levels such as for example to take the control of driving when mediation is necessary, making driving related decisions or for the purpose of other tasks the user is engaged in (such as online meeting of collogues using the car's infrastructure). This often involves two types of interaction: interaction for the purpose of the primary task (what the user is currently doing, for example playing a game) and the interaction for secondary tasks while the user is engaged in and her attention is on another primary task (such as for example, for receiving the status of the traffic ahead while the user is playing a game with the aid of car's infrastructure). Towards better facilitating these interactions, autonomous vehicles can benefit from multimodal interaction, the use of different (and often multiple) human sensory modalities. One the one hand, not much research has been done on how different human sensory modalities can be best used to facilitate in-car interaction. On the other hand, users of autonomous vehicles, use-situations, and what interaction requirements these users have remain largely unexplored. Except for speech and haptic based interaction, a little is known about using other modalities such as gestures, olfaction and sonification and how they can be used for different types of users and use situations. We aim to fill this gap by exploring how various types of sensory modalities can be used to enrich in-car interaction of various user groups and use situations.

References

1 Ranasinghe C., Holländer K., Currano R., Sirkin D., Moore D., Ju W., Schneegass S. *Autonomous vehicle-pedestrian interaction across cultures: towards designing better external Human Machine Interfaces (eHMIs)*. Extended Abstracts of the 2020 CHI Conference on Human Factors in Computing Systems, Hawaii, USA, 2020.

2 Janssen, C.P. and Kun, A.L. *Automated driving: getting and keeping the human in the loop*. Interactions, 27(2), pp.62-65.

3.4 A Multi-level Approach to Understanding the Now and the Future of Work and Wellbeing

Anna Cox (University College London, GB)

The creators of digital technology promised that digital tools would increase our productivity and give us more autonomy over when and where we work. Instead, it often seems as though increased use of information technology has reduced worker productivity (the productivity paradox) and led to us working always and from everywhere (the autonomy paradox). One interesting question is, how to create digital information technologies that actually deliver on these early promises so as to support both our productivity and wellbeing.

In order to answer this question we need to understand the role that technology plays in shaping how we work at multiple levels. The micro or task level enables us to understand how technologies support us in doing a particular task. Examples include studies of which emails people prioritise answering and why [1], and how can we design interventions to help people keep their focus on their work and better self-regulate self-interruptions [2]. Investigations at the meso or job level involve explorations of the influence of technology on the shape of a job or the development of relationships between team members. Examples include studies how people adapted to the switch to remote working brought on by the pandemic [3] and exploring how we can use videogames to create swift trust between remote teammates [4]. Explorations at the macro or life and wellbeing level enable us to focus on how technology influences the whole person, including outside of their work. Examples include studies of how people use technology to manage work life boundaries [5] and how we can use video games to support post-work recovery [6].

References

1 Anna L. Cox, Jon Bird, Duncan P. Brumby, Marta E. Cecchinato & Sandy J. J. Gould (2020) Prioritizing unread e-mails: people send urgent responses before important or short ones, Human–Computer Interaction, DOI: 10.1080/07370024.2020.1835481
2 Judith Borghouts, Duncan P. Brumby, and Anna L. Cox. 2020. TimeToFocus: Feedback on Interruption Durations Discourages Distractions and Shortens Interruptions. ACM Trans. Comput.-Hum. Interact. 27, 5, Article 32 (October 2020), 31 pages. DOI: 10.1145/3396044
3 Rudnicka, A., Newbold, J. W., Cook, D., Cecchinato, M. E., Gould, S., & Cox, A. L. (2020, August). Eworklife: Developing effective strategies for remote working during the COVID-19 pandemic. In The New Future of Work Online Symposium.
4 Evelyn Tan and Anna L. Cox. 2019. Trusted Teammates: Commercial Digital Games Can Be Effective Trust-Building Tools. CHI PLAY '19 Extended Abstracts. Association for Computing Machinery, New York, NY, USA, 705–713. DOI: 10.1145/3341215.3356296
5 Marta E. Cecchinato, Anna L. Cox, and Jon Bird. 2015. Working 9-5?: Professional Differences in Email and Boundary Management Practices In Proceedings of the 33rd Annual ACM Conference on Human Factors in Computing Systems (CHI '15). Association for Computing Machinery, New York, NY, USA, 3989–3998. DOI:10.1145/2702123.2702537
6 Collins, E., Cox, A., Wilcock, C., & Sethu-Jones, G. (2019). Digital games and mindfulness apps: comparison of effects on post work recovery. JMIR mental health, 6(7), e12853.

3.5 Designing to Support People Working, Connecting and Living Well in Mobile Work Environments

Geraldine Fitzpatrick (TU Wien, AT)

Technology is making it increasingly easier to connect, interact and work independent of traditional office spaces – work is becoming more mobile, more distributed in space and time. With this mobility comes a number of challenges. One is around how to be effective at/in work, to maintain motivation and productivity without the structures and oversight of work places. This requires people to operate with a high degree of autonomy and self-efficacy. And it requires new forms of leadership to enable such autonomy, provide appropriate support and to build trust. A second challenge is how to build and maintain high quality social connections in the absence of contemporaneous co-location. This is both for relationship building and for effective collaboration, creativity and problem solving. Strong social emotional skills are required to build and maintain relationships online, to create empathic connections, and to communicate effectively, often having to take more proactive and explicit steps for communications and interactions that could happen much more implicitly and serendipitously when co-located. A third challenge is about how to navigate time and space and work and all other aspects of life, often talked about in terms of blurring of boundaries. This can have both positive implications for increased flexibility and autonomy, mentioned above. It can also have negative implications for increasing stress and decreasing mental and physical health and well-being.

In our human-centred research, we explore roles for technology to support reflective work practices, e.g., [3], to develop emotional and social skills, e.g., [2], and to promote mental and physical health and well-being, e.g., [1].

References

1 Michel, T., Tachtler, F., Slovak, P. Fitzpatrick, G. (2020) Young People's Attitude Toward Positive Psychology Interventions: Thematic Analysis. JMIR Human Factors 7(4):e21145 2020, DOI: 10.2196/21145

2 Murphy, D, Slovak, P., Thieme, A., Jackson, D., Olivier, P., Fitzpatrick, G. (2019) Developing technology to enhance learning interpersonal skills in counselor education. British Journal of Guidance and Counselling, 47(3), 2019, pp. 328-341. http://dx.doi.org/10.1080/03069885.2017.1377337

3 Slovak, P., Frauenberger, C. and Fitzpatrick, G. Reflective Practicum: a Framework of Sensitising Concepts to Design for Transformative Reflection. In Proc CHI 2017, ACM, New York, NY, pp. 2696-2707. DOI: https://doi.org/10.1145/3025453.3025516

3.6 Is the Future of Work on the Move?

Christian P. Janssen (Utrecht University, NL)

Main reference Christian P. Janssen, Shamsi T. Iqbal, Andrew L. Kun, Stella F. Donker: "Interrupted by my car? Implications of interruption and interleaving research for automated vehicles", Int. J. Hum. Comput. Stud., Vol. 130, pp. 221–233, 2019.
URL http://dx.doi.org/10.1016/j.ijhcs.2019.07.004

Automated technology is changing human lives rapidly [1], including in the automotive domain. As vehicles get equipped with more and more reliable automated technology [2], there might be occasional times where the human driver does not need to pay full attention to the road and can or wants to, temporarily, pay attention to other tasks and activities. Doing other activities during an automated drive is a desire for many humans [3]. However, there is a potential of being distracted for too long and missing critical events or failing to respond timely to an in-car alert.

In our previous work, we have proposed that it can be beneficial to think of such scenarios in terms of interruption management: how do people interleave their attention between driving and non-driving activities [4]? The general idea is that people might not always respond immediately to an alert, but respond more gradually, consistent with how they interleave their attention in other multitasking settings [5] and consistent with the idea that auditory alerts might not always be processed when people are distracted [6, 7]. We have started to test these ideas experimentally, and so far see that drivers might indeed follow interleaving patterns when they are interrupted by an alert during semi-automated driving conditions [8]. This work shows the value that theories [4] and models [5, 9, 10] of human behavior can have: they can predict human behavior and guide the design of future interfaces.

One open question is how human behavior changes over time. As humans get experience with new interfaces (e.g., novel in-car technology), and are exposed to different environments and settings (e.g., different levels or forms of automation), their behavior might be different from what science might have predicted so far. Nonetheless, humans are humans, and theory about human behavior can guide and inform such insights. Another way that the future of work might be on the move more radically is whether humans will even travel as much by car as some did so far. Certainly, the recent pandemic has opened the eyes of many to the options of working from home. Again, thinking about such environments from an interruption perspective is beneficial, as theories of interruption have been proposed in both automotive [4] as well as office and home settings [5]. Although this might create some limitations (like being distracted by the home environment and pets and family members), there are also opportunities. For example, on a personal note, I was able to attend this Dagstuhl seminar remotely, despite having become a father recently. I am sure that my own future of work will change quite a bit over the next few years.

References

1. Janssen, C. P., Donker, S. F., Brumby, D. P., & Kun, A. L. (2019). History and future of human-automation interaction. International Journal of Human-Computer Studies, 131, 99-107.
2. Kun, A. L. (2018). Human-machine interaction for vehicles: Review and outlook. Foundations and Trends in Human–Computer Interaction, 11(4), 201-293.
3. Pfleging, B., Rang, M., & Broy, N. (2016, December). Investigating user needs for non-driving-related activities during automated driving. In Proceedings of the 15th international conference on mobile and ubiquitous multimedia (pp. 91-99).
4. Janssen, C. P., Iqbal, S. T., Kun, A. L., & Donker, S. F. (2019). Interrupted by my car? Implications of interruption and interleaving research for automated vehicles. International Journal of Human-Computer Studies, 130, 221-233.
5. Brumby, D. P., Janssen, C. P., & Mark, G. (2019). How do interruptions affect productivity?. In Rethinking Productivity in Software Engineering (pp. 85-107). Apress, Berkeley, CA.
6. Van der Heiden, R. M., Janssen, C. P., Donker, S. F., Hardeman, L. E., Mans, K., & Kenemans, J. L. (2018). Susceptibility to audio signals during autonomous driving. PloS one, 13(8), e0201963.
7. Van der Heiden, R. M., Kenemans, J. L., Donker, S. F., & Janssen, C. P. (2021, online first). The effect of cognitive load on auditory susceptibility during automated driving. Human factors, https://doi.org/10.1177/0018720821998850
8. Nagaraju, D., Ansah, A., A., Ch, N.A.N., Mills, C., Janssen, C.P., Shaer, O., and Kun, A.L. (2021) How Will Drivers Take Back Control in Automated Vehicles? Driving Simulator Test of and Interleaving Framework. In: Proceedings of the ACM Conference on Automotive User Interfaces and Interactive Vehicular Applications. New York, NY: ACM Press. https://doi.org/10.1145/1122445.1122456
9. Brumby, D. P., Janssen, C. P., Kujala, T., & Salvucci, D. D. (2018). Computational models of user multitasking. In: Computational interaction design, 341-362.
10. Janssen, C. P., Boyle, L. N., Ju, W., Riener, A., & Alvarez, I. (2020). Agents, environments, scenarios: A framework for examining models and simulations of human-vehicle interaction. Transportation research interdisciplinary perspectives, 8, 100214.

3.7 The Future of Work and Wellbeing: From Automated Vehicles to Working at Home

Andrew Kun (University of New Hampshire – Durham, US)

Automated vehicles hold the promise of allowing us to use the time we spend on the road for productive endeavors, or for attending to our wellbeing. Thus, future commutes to work might include preparation for the workday. Similarly, trips from the office back home might allow us the time and place to shift out of work-mode and into home-life-mode. One interesting question is, how to create user interfaces that will allow us to take advantage of automated vehicles such as to support our productivity and wellbeing. We need to explore this question in light of the fact that automated vehicles are likely to be only partially automated for many years to come. This means that driving will be part of what we do on a trip in a vehicle, and user interface design will have to support driving safety [1]; however there will be extended periods of time (e.g. on a stretch of highway), where automation can take over, and we can focus on non-driving tasks. Periodically, automation will interrupt us in the non-driving task, and we will need to take back control of the vehicle [2].

Of course, in the coming years many workers might commute significantly less than in the past [3]. Yet, as user interface designers, we can notice parallels between the engagement in work tasks in a vehicle, and working from home. In both cases we work in an environment that possibly lacks all of the tools that are in the office or our lab. We are also interrupted unexpectedly in both environments – in the vehicle the interruption comes from the automation that needs us to take over control, while at home it might be the needs of a child. In both cases our co-workers are remote, and our communication with them has to happen in the context of distractions and interruptions. Thus, as we design interfaces for automated vehicles and for working from home, it will be fruitful to understand the similarities of the two domains.

References

1 Andrew L. Kun. Human-machine interaction for vehicles: Review and outlook. Foundations and Trends in Human–Computer Interaction 11, no. 4 (2018): 201-293.

2 Christian P. Janssen, Shamsi T. Iqbal, Andrew L. Kun, and Stella F. Donker. Interrupted by my car? Implications of interruption and interleaving research for automated vehicles. International Journal of Human-Computer Studies 130 (2019): 221-233.

3 A. Kun, R. Sadun, O. Shaer, and T. Teodorovicz. Where did the commute time go. Harvard Business Review. Available online: https://hbr. org/2020/12/where-did-the-commute-time-go (accessed on 3 January 2021) (2020).

3.8 The Impact of XR-Headsets on Mobile Productivity

Mark McGill (University of Glasgow, GB)

Main reference Mark McGill, Aidan Kehoe, Euan Freeman, Stephen A. Brewster: "Expanding the Bounds of Seated Virtual Workspaces", ACM Trans. Comput. Hum. Interact., Vol. 27(3), pp. 13:1–13:40, 2020.
URL http://dx.doi.org/10.1145/3380959

As part of ERC ViAjeRo (`https://viajero-project.org/`), my research has focused on some of the key benefits and challenges offered by adopting XR headsets (augmented/virtual reality) in passenger contexts, offering new avenues for productivity and entertainment whilst potentially overcoming key impediments such as motion sickness. With respect to productivity, there are some notable benefits offered by transitioning from physical displays to virtual content and displays rendered by XR headsets. For example in terms of ergonomics we can dynamically manipulate the virtual content position for comfort and accessibility [1], making the most of our restricted capacity for movement in constrained spaces [2] like an economy plane seat [3].

However, we're also finding that a transition toward XR content in passenger contexts is posing unique challenges beyond the technical, for example in terms of social acceptability constraining how users see themselves adopting these technologies around other passengers. I believe that if we can resolve these key challenges in difficult mobility contexts, those solutions will also positively impact everyday productivity in many more contexts e.g. working from home, helping to adapt the home environment toward productivity just as we adapt the car or plane interior.

References

1 Mark Mcgill, Aidan Kehoe, Euan Freeman, and Stephen Brewster. 2020. Expanding the Bounds of Seated Virtual Workspaces. ACM Trans. Comput.-Hum. Interact. 27, 3, Article 13 (June 2020), 40 pages. DOI: https://doi.org/10.1145/3380959
2 McGill, M., Williamson, J., Ng, A. et al. Challenges in passenger use of mixed reality headsets in cars and other transportation. Virtual Reality 24, 583–603 (2020). https://doi.org/10.1007/s10055-019-00420-x
3 Julie R. Williamson, Mark McGill, and Khari Outram. 2019. PlaneVR: Social Acceptability of Virtual Reality for Aeroplane Passengers. Proceedings of the 2019 CHI Conference on Human Factors in Computing Systems. Association for Computing Machinery, New York, NY, USA, Paper 80, 1–14. DOI:https://doi.org/10.1145/3290605.3300310

3.9 Challenges in Mobile Offices in Automated Vehicles

Alexander Meschtscherjakov (Universität Salzburg, AT)

In the near future more and more Advanced Driving Assistance Systems (ADAS) will be integrated into modern vehicles and cars will be able to drive at different levels of automation. This relieves drivers from driving tasks allowing them to be involved in other activities – may it be work or leisure related. Especially for commuting this fact offers potentials to conduct work already while commuting to and from work. Depending on the type of work I see the following challenges to be addressed, to make mobile commuting work a success:

1. Allocation of driving related responsibilities between the automated vehicle and the driver needs to be defined and communicated
2. Automation mode must ensure a certain time frame of reliving drivers from any driving task related activities (e.g., monitoring the environment, take-over requests below a reasonable time span)
3. Providing infrastructure that allows mobile working while riding in a vehicle
4. Addressing aspects of motion sickness, re-routing, etc.

3.10 Resolving of Automation Level Issues and Operational Constraints to Enable Productive Work in AVs

Andreas Riener (TH Ingolstadt, DE)

Main reference Andreas Riener, Myounghoon Jeon, Ignacio Alvarez (eds).: "User Experience Design in the Era of Automated Driving", Springer Studies in Computational Intelligence, pp. 623, 2021
URL https://doi.org/10.1007/978-3-030-77726-5

Automated driving has undergone a major development boost in recent years, but the initial hype has also suffered some setbacks in the form of reality checks. While Level 3 automation was expected to hit the road by 2021 on a broad scale, we see today that this is no longer the case, mainly because of safety and legal issues [1]. In addition, recent accidents involving automated vehicles have fueled public fears, both on the side of policy makers and the general society. Legal and regulatory frameworks are subject to frequent changes and further vary from country to country. While there are many unanswered questions and uncertainties surrounding automated driving, broad user acceptance is considered a basic prerequisite for using these systems and thus also for bringing applications for productive work, among other things, into the vehicle. One particular issue discussed in this context is in which automation level effective working will be possible. There is broad consensus that this should be the case on levels 3 [2] or higher, because below that the driver is obliged to permanently monitor the vehicle. A common misconception is, however, that a specific car is a level 'X' automated vehicle – this is not correct! A specific vehicle might be equipped with several advanced driver assistance systems/driving automation systems implemented at different levels of automation. And for a specific function, its availability is further restricted by the constraints of the "Operational Design Domain" (ODD). ODD specifies the boundary conditions for a specific automation function, including environmental, geographical, time-of-day restrictions, and/or the requisite presence or absence of certain traffic or roadway characteristics for the intended automation functionality. For example, a (L3) highway assist might only operate on the (German) highway A9 from Munich to Berlin, under clear weather, dry road, visible lane markings, in the speed range 45 to 130km/h, and out of rush hour. Outside these boundary conditions, the system is not available (and may switch off automatically). The combination of assistance systems or automated driving functions on different levels in a variety of application scenarios (and restrictions) leads to ambiguity and high uncertainty and makes it almost impossible for the driver/passenger to find out whether an assistance or automated driving system is available in a certain situation, which one, and at which level of automation. Also, the availability might change rather quickly, based on the definition of the ODD. This example shows how difficult it is to implement efficient mobile working systems for the automated vehicle while ensuring that driving safety is not compromised by its use.

The consequence is either systems that have to adapt frequently and quickly to the current situation (level, ODD), which, depending on the context, has a negative impact on efficiency of work, or to wait with its introduction until at least level 4 is available on a large part of the commuter route. In our research group at Technische Hochschule Ingolstadt my PhD students and postdoctoral researchers explore together with me different opportunities to support the driver-passenger of an automated vehicle, for example, by implementing (personalizable, adaptive) interfaces to support mode awareness [3], develop and test prototypes for efficient and ergonomic office work in Level 3+ vehicles [4], build (mobile) driving simulators to test our interfaces with the general public, study the capability of augmented and/or virtual reality as an additional level of interaction [5], create (transparency) displays to communicate upcoming decisions/maneuvers from the vehicle to the passenger, investigate the potential that driver-automation cooperation has on road safety, to name a few.

I am looking forward to the seminar in Dagstuhl (although it is virtual this time) to discuss these challenges together, to identify/adapt solutions (possibly also from other domains) and to create a roadmap of research with all the participating experts.

References

1 Andreas Riener, Myounghoon Jeon, and Ignacio Alvarez (eds). “User Experience Design in the Era of Automated Driving”, Springer Studies in Computational Intelligence, pp. 623, 2021, ISBN: 978-3-030-77725-8, DOI: 10.1007/978-3-030-77726-5

2 SAE J3016-2021 – Taxonomy and Definitions for Terms Related to Driving Automation Systems for On-Road Motor Vehicles. SAE International, (April 2021).

3 Philipp Wintersberger, Hannah Nicklas, Thomas Martlbauer, Stephan Hammer, and Andreas Riener. 2020. Explainable Automation: Personalized and Adaptive UIs to Foster Trust and Understanding of Driving Automation Systems. In 12th International Conference on Automotive User Interfaces and Interactive Vehicular Applications (AutomotiveUI ’20). Association for Computing Machinery, New York, NY, USA, 252–261. DOI: https://doi.org/10.1145/3409120.3410659

4 Clemens Schartmüller, Klemens Weigl, Andreas Löcken, Philipp Wintersberger, Marco Steinhauser, and Andreas Riener. 2021. “Displays for Productive Non-Driving Related Tasks: Visual Behavior and Its Impact in Conditionally Automated Driving” Multimodal Technologies and Interaction 5, no. 4: 21. https://doi.org/10.3390/mti5040021

5 Andreas Riegler, Andreas Riener, and Clemens Holzmann. 2020. A Research Agenda for Mixed Reality in Automated Vehicles. In 19th International Conference on Mobile and Ubiquitous Multimedia (MUM 2020). Association for Computing Machinery, New York, NY, USA, 119–131. DOI: https://doi.org/10.1145/3428361.3428390

3.11 Human-Computer Interaction to Support Work and Wellbeing in Mobile Environments

Sayan Sarcar (University of Tsukuba – Ibaraki, JP)

Of late, as the ubiquitous computing field progresses, mobile and wearable devices have become mainstream computing resources to be operated anywhere and anytime. However, it has been an alarming fact that excessive use of mobile devices imposes a negative effect on the health and wellbeing of users. As Human-computer Interaction (HCI) researchers, our challenge is to design effective ways to use technology in a balanced way between work and wellbeing.

In this online Dagstuhl Seminar, HCI experts discussed possible research questions pertinent to this underexplored research area as follows.

- What are the manual, visual, auditory, and cognitive demands of tasks in mobile environments? Many present and future scenarios were discussed where such demands are prominent, such as a mobile office in the automated car environments, virtual environments such as AR, VR, MR, XR spaces for game playing, or doing office works. Also, the task spaces in such environments were explored with possible ways to make them less cognitively demanding.
- What are the aspects of mobile contexts that affect how people work and play in mobile environments? The workshop attendees also discussed several aspects: distraction caused by social media, less focus on the work users currently doing; scarcity of programs that can dynamically manipulate the work-life imbalances and encourage users to perform tasks to help balancing it.
- How do we support (safe) task switching?

Some potential strategies were discussed – such as how to incorporate programs to encourage short-term physical or mental exercising between two tasks to improve the focus, natural ways to switch the tasks.

How do we leverage advanced HCI technologies to support work and wellbeing activities in mobile environments? The discussion was focused around introducing immersive technologies (AR, VR, MR, XR) and design supportive interventions.

3.12 Productive and Safe UIs in Conditionally Automated Driving

Clemens Schartmüller (TH Ingolstadt, DE)

Main reference Clemens Schartmüller, Philipp Wintersberger, Anna-Katharina Frison, Andreas Riener: "Type-o-Steer: Reimagining the Steering Wheel for Productive Non-Driving Related Tasks in Conditionally Automated Vehicles", in Proc. of the 2019 IEEE Intelligent Vehicles Symposium, IV 2019, Paris, France, June 9-12, 2019, pp. 1699–1706, IEEE, 2019.
URL http://dx.doi.org/10.1109/IVS.2019.8814088

The COVID-19 pandemic has shown how technology enables knowledge workers to work literally everywhere. However, we have also seen how important personal contacts are to us as social human beings, also in the context of work. While traveling for personal contacts was previously often seen as a waste of time and resources, future automated vehicles may enable us to use this time productively. Furthermore, it may even allow retaining, e.g., the jobs of truck drivers. In a mobile office automated vehicle, they can do logistics planning between on- and offloading, where they would otherwise be rationalized away by driving automation.

However, especially the first levels of automated driving will inherit safety-critical transitions of control from automated driving to manual driving. Accordingly, user interfaces are needed that enable safe working in such vehicles. We have quantitatively shown that merging knowledge from office ergonomics and human factors in automated driving can result in safe and productive in-vehicle user interfaces [1]. However, open questions include how novel technology like mixed reality can further foster safety and productivity without confusing the user [2] – do we need to deviate from the old interaction paradigms like WIMP wimp? – and, in the grand scheme of things, how will the ability to work on the go impact our daily routines, will there be a place left for "shutting down" or is a mobile office just another contributor to the "always-on" pressure?

References

1 Schartmüller, C., Wintersberger, P., Frison, A.-K., & Riener, A. (2019). Type-o-Steer: Reimagining the Steering Wheel for Productive Non-Driving Related Tasks in Conditionally Automated Vehicles. Proceedings of the 30th IEEE Intelligent Vehicles Symposium, IV'19, 1489–1496. https://doi.org/10.1109/IVS.2019.8814088

2 Schartmüller, C., Riener, A., Wintersberger, P., & Frison, A.-K. (2018). Workaholistic: On Balancing Typing- and Handover-Performance in Automated Driving. Proceedings of the 20th International Conference on Human-Computer Interaction with Mobile Devices and Services, MobileHCI'18, 16:1–16:12. https://doi.org/10.1145/3229434.3229459

3 Schartmüller, C., & Riener, A. (2018). Safe, Productive and Socially Accepted Text Input in Highly Automated Driving. MobileHCI 2018 Workshop on Socio-Technical Aspects of Text Entry, 18–23. http://ceur-ws.org/Vol-2183/position3.pdf

3.13 The Future of Work and Mobility in Automated Vehicles

Martina Schuß (TH Ingolstadt, DE)

Main reference Martina Schuß, Philipp Wintersberger, Andreas Riener: "Let's Share a Ride into the Future: A Qualitative Study Comparing Hypothetical Implementation Scenarios of Automated Vehicles", in Proc. of the 2021 CHI Conference on Human Factors in Computing Systems, CHI '21, Association for Computing Machinery, 2021.
URL http://dx.doi.org/10.1145/3411764.3445609

In the future automated vehicles (AVs) will permanently change our mobility. At the moment it is uncertain how these vehicles will be implemented and several scenarios including private car ownership, as well as scenarios that require sharing of rides and vehicles are possible. Most probably, shared automated vehicles (SAVs) where many rides are shared among people will be the most beneficial ones for society and environment, and they are my research focus. However, research in this area often focuses on the usability of HMI concepts rather than on user's emotions or daily life context, and includes unbalanced participant groups in terms of gender, age, and other social or psychological identities. In my research I am focusing on shared automated vehicles and I am taking a pluralistic viewpoint paying attention to groups of people that are many times are left out in the research and development of new technologies, such as people from different cultural backgrounds, senior citizens but also women and children [1].

During the Dagstuhl workshop "Workshop for the Future of Work and Mobility in Automated Vehicles" that was hosted on 8th June 2021 (see below for a more detailed summary), participants were vividly discussing the fulfillment of needs in SAVs in the context of work. Some of the most interesting findings were that autonomy and security were the least fulfilled needs in shared contexts and especially concerns due to co-passengers are unexpectedly complex and important. In the future, concepts are needed to account for a high feeling of autonomy and security in SAVs to make them the mobility of choice for as many people as possible.

References

1 M. Schuß, P. Wintersberger, and A. Riener. 2021. "Lets Share a Ride into the Future: A Qualitative Study Comparing Hypothetical Implementation Scenarios of Automated Vehicles." Proceedings of the 2021 CHI Conference on Human Factors in Computing Systems.

3.14 Existential Embodied Agents for Home and Mobile Environments

Gregory F. Welch (University of Central Florida – Orlando, US)

Main reference Kangsoo Kim, Luke Boelling, Steffen Haesler, Jeremy N. Bailenson, Gerd Bruder, Gregory F. Welch: "Does a Digital Assistant Need a Body? The Influence of Visual Embodiment and Social Behavior on the Perception of Intelligent Virtual Agents in AR", in Proc. of the IEEE International Symposium on Mixed and Augmented Reality, ISMAR 2018, Munich, Germany, October 16-20, 2018, pp. 105–114, IEEE, 2018.
URL http://dx.doi.org/10.1109/ISMAR.2018.00039

The future of work and wellbeing contextualized in mobile environments could actually be seen as yet another encroachment of work on life and wellbeing. Yet we (humans) are in control of the technology, and have the opportunity to make it work better for us for both work and wellbeing. Beyond direct interfaces to machines, autonomous agents have been increasing in acceptance and corresponding use. This is perhaps nowhere more eviden than with today's "smart speakers" from Amazon, Apple, and Google. While these devices are becoming more powerful, today's devices are inherently transactional, reactive, and relatively neutral with respect to influence. Beyond solely audio, embodied agents (agents with visual form) offer many more affordances in terms of their ability to communicate (adding nonverbal aspects) and their sense of existence. Whether in the home or in a mobile environment, such agents can provide autonomous intelligent behavior that is proactive, personalized, and adaptive to our needs and circumstances, while also providing a common "clean" interface to machines. Such agents can acton behalf of individuals and service providers, exerting influence aimed at short and long-term goals. Agents who are continuously existent, with an apparent independent virtual life, will be able to transform what is now a transactional relationship (e.g., "What time is it?") to a more relational one, with increased influence resulting from our perception of an existential companion.

References

1 Kangsoo Kim; Luke Boelling; Steffen Haesler; Jeremy N. Bailenson; Gerd Bruder; Gregory F. Welch. Does a Digital Assistant Need a Body? The Influence of Visual Embodiment and Social Behavior on the Perception of Intelligent Virtual Agents in AR. In: Proceedings of the 17th IEEE International Symposium on Mixed and Augmented Reality (ISMAR 2018), Munich, Germany, October 16–20, 2018, 2018.

4 Working groups

4.1 Reporting on Bilateral Group Activity: Curated Conversation

Sun Joo Ahn (University of Georgia – Athens, US) and Gregory F. Welch (University of Central Florida – Orlando, US)

URL https://youtu.be/UmEuKcwaJWU

On Thursday, June 10, Grace Ahn and Greg Welch held a conversation on the utility and perils of historical data collected by autonomous vehicles, in which we discussed the dilemma of the utility of accumulated data for autonomous vehicles versus the dangers posed in control, ownership, and management of the collected data. Autonomous vehicles (and related

AI agents) cannot optimize their utility for users without collecting personal and behavioral data. Users may recognize that their direct interaction data with autonomous vehicles may be tracked and logged but may not be aware that other personal data may also be collected as the autonomous vehicle "gets to know" its user, such as medical data, purchase data, daily schedules, and social network data related to the users' friends and family. Users will want to maximize the convenience and utility of autonomous vehicles, but how can they protect third parties from gaining control over the collected data? Third parties, such as insurance companies and employers, may be able to make unwarranted inferences from the data collected on these users. This can be dangerous when people have a positivity bias toward data collected and processed by machines, considering these results to be more accurate and "correct" than humans.

One potential solution may be to allow users to grant temporary access to a very specific "lot" of data to increase the utility of specific goods and services, and then be able to revoke access as well as delete data that were temporarily provided. This voluntary provision of data may also assist in developing and training new autonomous platforms. Earlier research suggests that users are not necessarily against providing personal data, as long as there is utility to be gained. Users would appreciate more visibility and transparency in how they are better able to manage and control the personal data being collected, and this in turn, is likely to encourage data sharing. Collection of personal data is typically seen as an "evil" tactic to manipulate users, but many of the services and conveniences we benefit from would not be possible without collection of personal data. Therefore, we should recognize that data collection is a highly nuanced and contextual activity, its utility and value resting on the entities involved and for what reason the data are being collected.

Link to YouTube video of the conversation: `https://youtu.be/UmEuKcwaJWU`

4.2 Reporting on Team Breakout Activity: Curated Conversation

Mark McGill (University of Glasgow, GB), Laura Boffi (University of Ferrara, IT), and Susanne Boll (Universität Oldenburg, DE)

URL https://www.youtube.com/watch?v=lV3AyLzKtS0

On Thursday 10th June, our breakout group (Susanne Boll, Laura Boffi, Mark McGill) had a conversation around the future of work and autonomous vehicles. Based on prior group activities earlier in the week, we chose three topics noted to have the most significant impact:

- What kind of experiences we want to facilitate in autonomous vehicles – This considered the core aim of the workshop, to facilitate work in mobile environments. Our discussion reflected on existing use of travel time, and the possibility that by operationalizing travel, we might lose opportunities for reflection.
- Motion sickness and the fixation on visual experiences – Given the challenge of resolving motion sickness, the group approached this problem from the perspective of ruling out visually-led AV experiences, discussing the capacity of auditory presence / auditory virtual reality to facilitate passenger presence in other auditory environments, as a means to escape the journey.
- Algorithmic bias in shared AVs – Inspired by the group movie activity, we finally discussed sources of algorithmic bias as might be exhibited in rideshare platforms, and the potential societal impact that such bias might have given a transition from owned to shared transit,

and a developed reliance on this as the primary means of public transport. For example, consider structural bias in what areas AVs travel to or through, or how they collect passengers in different areas, or based on different socioeconomic profiles. Such bias' suggest that we might reinforce existing inequalities in shared AV services, or even develop new ones based on the profitability of servicing certain locations.

See the main reference of this abstract for a recording of this discussion.

4.3 Digital Assistants for Cars (Youtube-Playlist)

Andreas Riener (TH Ingolstadt, DE) and Gary Burnett (University of Nottingham, GB)

URL https://www.youtube.com/playlist?list=PLJi9U1fnyFjdPVbtYmDfmfdE4mPIrfh1l

Below, Gary and I (see Fig. 10) share a selection of videos showing different concepts for digital assistants/Intelligent Personal Assistants (IPAs), user-friendly vehicle interfaces in infotainment system, or interactive human-machine devices that react to the user's emotional state. On the one hand, this allows users to be more actively involved in these vehicles, and on the other hand, it also provides customers with the opportunity for better cooperation and effective collaboration with the systems. Some of the videos selected here show real digital assistants, some represent visions, some are jokes (but still informative). They differ in the way of representation (avatar, other form of visual embodiment, features of anthropomorphism in visualization, voice, natural language, etc.).

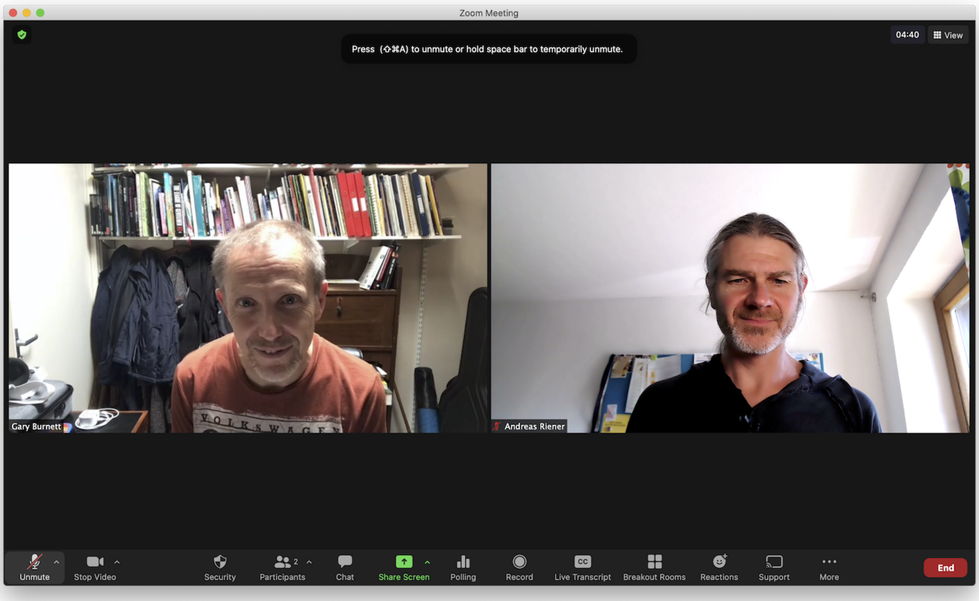

Figure 10 Gary Burnett and Andreas Riener during the Thursday group activity: Compiling a Youtube playlist.

4.3.1 Youtube playlist

The Youtube playlist with our selection can be watched here: `https://www.youtube.com/playlist?list=PLJi9U1fnyFjdPVbtYmDfmfdE4mPIrfh1l`.

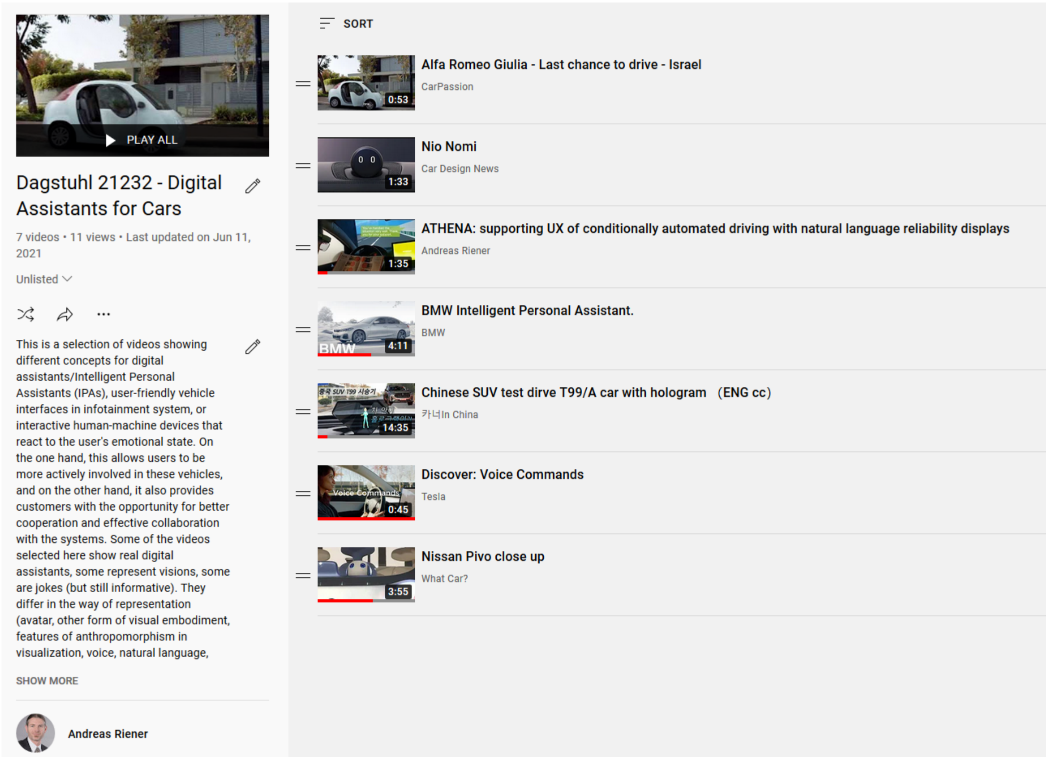

Figure 11 Youtube playlist with our selection of "must-seen" videos.

In summary, we can say that digital (or personal) assistants have great potential to support interaction, communication (incl. work) in the vehicle in the future.

4.3.2 Extended, full list of inspiring videos

- A joke advert but demonstrating many of the things you don't want to do with a digital assistant design in an automated car, `https://www.youtube.com/watch?v=UIzwQbOOBkY`
- Faw Bestune Chinese SUV with dancing hologram avatar on the dashboard (demonstrates many potential EDI issues), `https://www.youtube.com/watch?v=XQnvSKLLPCO`
- AIDA 1.0 – MIT SENSEable City Lab and Personal Robots Group with collaboration from Audi VW (video from 2014 but system created in 2009). A robot-like digital assistant, `https://www.youtube.com/watch?v=jCiTYytpMpQ`
- AIDA 2.0 – MIT SENSEable City Lab with collaboration from Audi VW (2011). They removed the robot-like presence, `https://www.youtube.com/watch?v=V9Qmg4TteMY`
- BMW's Intelligent Personal Assistant (not robotic/avatar based but it's what we currently can see/use), `https://youtu.be/C-gRJrDrICs`, `https://youtu.be/NP-ZzuKAD8k`
- Tesla's video on using voice commands to interact with their cars (I've never driven a Tesla, but it seems they take a "less communicative" approach unless you have a command/request for the car and most actions can be triggered using the screen or using speech), `https://youtu.be/oDruklAJFmA`

- Nissan's PIVO 2 (2007 concept car), https://www.youtube.com/watch?v=Lb1UYLhrFhw
- Penguin-like in0car robot by Pioneer, https://www.youtube.com/watch?v=GlSXHKJ6MDQ
- Another oldish concept by BYD that never made it to the roads, https://www.youtube.com/watch?v=BqncDHOmPlM

4.4 Reporting on Bilateral Group Activity: Curated Conversation about "Don't forget about physicality in HCI"

Clemens Schartmüller (TH Ingolstadt, DE)

Joint work of Clemens Schartmüller, Champika Manuel Epa Ranasinghe
URL https://youtu.be/zQ80SNHHDIY

In the day 4 bilateral activity, we, Champika Manuel Epa Ranasinghe and Clemens Schartmüller discussed how physicality is often overlooked in recent HCI research. The discussion was recorded on video, available here: https://youtu.be/zQ80SNHHDIY.

First, we discussed how physical interaction has the potential to improve the memorability of interactions, provide interactivity to disadvantaged groups like elderly and disabled people, as well as improve overall well-being. However, current mobile devices and also vehicles are trending toward touch displays, which not only results in a loss of the full haptic sensation of pressing a button but were also shown to be detrimental to driver safety. We then found that digital techniques and physical interfaces are easily combinable e.g., by combining RFID tags and tangibles. Coming back to work in automated vehicles, we discussed how slowing down with physical interfaces could be beneficial for users. Slowing down could support in being able to grasp the heavy information load presented in a mobile office (office tasks but also driving-related information) and thereby avoid users being just a passive recipient of information that misses half – after all, understanding, not only speed, is critical for effectiveness.

4.5 Reporting on the "Workshop for the Future of Work and Mobility in Automated Vehicles" Workshop

Clemens Schartmüller (TH Ingolstadt, DE), Andreas Riener (TH Ingolstadt, DE), and Martina Schuß (TH Ingolstadt, DE)

On Tuesday, June 8, the Dagstuhl seminar hosted a workshop entitled "Workshop for the Future of Work and Mobility in Automated Vehicles". Two researchers from Prof. Dr. Andreas Riener's HCI Group- Clemens Schartmüller and Martina Schuß- conducted the workshops during two different time slots. The workshop started with an introduction on automated driving in general and two forms of automated vehicles were presented to the workshop participants: private automated vehicles (PAVs) and shared automated vehicles (SAVs). Subsequently, the so-called user need cards [1-3] were presented to participants. These cards represent psychological needs that in situations or in the interaction with products can either be fulfilled or not fulfilled leading to either positive or negative user experiences. Thus, these needs should be taken into account when designing interfaces in the context of AD and interfaces should be designed to enhance user's needs. The workshop participants were then split up into two breakout rooms and were brainstorming on whether the user needs were fulfilled, respectively not fulfilled, in the context of work. Thus, one breakout room was discussing on PAVs while the other participants were discussing on SAVs. After the first discussion round and a short coffee break, the groups were ideating on possible design solutions on the least fulfilled needs in the context of work in either PAVs or SAVs. The workshop closed with a plenary wrap-up round and goodbye.

Below we briefly sum up which needs were the least fulfilled in the respective scenarios including AVs:

Fulfilled Not fulfilled PAVs Autonomy, popularity, security, Meaning, relatedness, physicalness **SAVs** Competence, meaning, relatedness Autonomy, popularity, security

Relatedness, physicalness, competence, meaning Security, relatedness, popularity, autonomy

References

1. Hassenzahl, M. & Diefenbach, S. (2012). Well-being, need fulfillment, and Experience Design. In Proceedings of the DIS 2012 Workshop on Designing Wellbeing. June 11-12, 2012, Newcastle, UK.
2. Hassenzahl, M. , Diefenbach, S. & Göritz, A. (2010). Needs, affect, and interactive products – Facets of user experience. Interacting with Computers, 22, 353-362.
3. Hassenzahl, M., Eckoldt, K., Diefenbach, S., Laschke, M. & Lenz, E. (2013). Designing moments of meaning and pleasure – Experience Design. International Journal of Design, 7 (3), 21-31.

5 Panel discussions

5.1 Report on "The Future of Creative Teams" panel

Andrew Kun (University of New Hampshire – Durham, US) and Orit Shaer (Wellesley College, US)

On Friday, June 11, the Dagstuhl seminar hosted a panel entitled "The future of creative teams." The panelists were two distinguished researchers, each of whom explores creative teamwork, and each of whom focuses on a different set of questions related to creativity and the future of (team) work.

The panel started with the two panelists each responding to the following question: "What is the future of creative teams through the lens of HCI?" After the panelists provided their responses, we proceeded with questions from the participants of the workshop, in which panelists elaborated on their main points. Below, we briefly introduce the two panelists and the central message they conveyed in the panel.

Panelists:

1. Mark Gross
   - Bio: Professor, University of Colorado, Boulder. Director of ATLAS Institutie. He is an expert on robotics, design, and tangible interaction.
   - Message: Creativity craves constraint. Diverse teams yield better outcomes. Focus on fun first.

2. Amon Millner
   - Bio: Associate Professor, Olin College. His research focuses on extending access to STEM empowerment.
   - Message: We must understand pre-existing and future social structures that surround how creative teams are formed, and influence (or even determine) the tools they use to carry out their work. How can we leverage all that we learned during the pandemic disruption to chart more equitable creative spaces and practices moving forward?

5.2 Report on the "Online Platforms and the Future of Work" panel

Andrew Kun (University of New Hampshire – Durham, US) and Orit Shaer (Wellesley College, US)

On Friday, June 11, the Dagstuhl seminar hosted a panel entitled "Online Platforms and the Future of Work." The panelists were three distinguished researchers, each of whom explores online platforms from the perspective of work, and each of whom focuses on a different set of questions related to these platforms.

The panel started with the three panelists each responding to the following question: "How will we, and/or how should we, use online platforms for the future of work?" After each of the three panelists provided their response, we proceeded with questions from the participants of the workshop, in which panelists elaborated on their main points. Below, we briefly introduce the three panelists and the central message they conveyed in the panel.

Panelists:

1. Jennifer Golbeck
   - Bio: Professor, University of Maryland, College Park. She is an expert in social networks, social media, privacy, and security on the web. She is also a content creator on various online platforms including Twitter and TikTok.
   - Message: Building an effective, engaged presence online – especially on social media – is a lot of work that's often dismissed. We need to value the content creation process more.
2. Neha Kumar
   - Bio: Associate Professor, Georgia Tech. Her research is at the intersection of human-centered computing and global development, and she works on technology design for/with communities that have historically been underserved.
   - Message: How do we ensure equality when workers use online platforms? Who are we leaving out and who are we keeping in, in the design of these online platforms?
3. Saiph Savage
   - Bio: Assistant Professor, Northeastern University. She is an expert on online labor platforms.
   - Message: We need auditing mechanisms to be able to understand what is happening inside online labor platforms. This is crucial to push for change (we need to understand problems that exist and also tways to put pressure on online labor platforms.)

6 Open problems

6.1 Summary and Next Steps

Andreas Riener (TH Ingolstadt, DE), Stephen Brewster (University of Glasgow, GB), Andrew Kun (University of New Hampshire – Durham, US), and Orit Shaer (Wellesley College, US)

With this Dagstuhl seminar, we set out to accomplish three goals. Here we review these goals and report on our progress towards accomplishing our original goals.

6.1.1 Overview of state-of-the-art technologies, methods, and models

The workshop was highly successful in providing both breadth and depth to attendees in reviewing the state-of-the-art in HCI for the future of work in mobile environments. Participants provided reports on their own work, and additionally five world-class experts from our community provided information on where technologies, methods and models are today.

6.1.2 List of challenges and hypotheses

Workshop participants identified several overarching challenges, as well as a long list of specific challenges. The overarching challenges centered around three issues. First, we focused on the algorithmic bias in how tools are designed for the future of work. These were underscored both from the discussions in our own meetings, and as we viewed the "Coded Bias" documentary. Second, we also focused on the global inequity in access to the knowledge and tools that will create and enable the future of work. Third, we discussed the lack of understanding how future work arrangements can help and hurt worker productivity, creativity, and very importantly, their wellbeing.

The list of specific problems included a host of problems in designing a broad swath of human-computer interfaces, not having information on how online platforms monitor, reward, and punish gig workers, and not having sufficient understanding of how the sudden change in work that was imposed by the COVID-19 pandemic affected workers, and what the long-term effects of this change will be.

We also worked towards formulating hypotheses to start the work of tackling the above challenges. Our work included both discussions about very specific problems, such as how to support a particular important activity (such as work in future automated vehicles), to how we can contribute to progress on overarching problems. The hypotheses that we discussed were focused on specific problems. One example is the suggestion of Saiph Savage that we need audit mechanisms for online platforms to effectively support gig workers. Another example is that, to support workers in future automated vehicles, we need to introduce user interfaces based on speech interaction, augmented reality, and tangible interfaces.

6.1.3 Roadmap(s) for research

Throughout this seminar we kept asking "which way do we move forward, together?" One avenue that we will pursue is to maintain the cohesiveness of the group that was assembled (virtually) for this Dagstuhl seminar by engagement in follow-on activities, such as the CHIWORK 2021-2022 symposium (www.chiwork.org). This symposium will feature weekly conversations on the topic of HCI and the future of work and wellbeing. We plan to continue our Dagstuhl conversations at this symposium.

We are also continuing to pursue international scientific collaboration that was supported and inspired by this Dagstuhl seminar. Along these lines, the two US-based organizers, Shaer and Kun, submitted a grant request to fund research experiences for US students to visit Germany. Two co-organizers (Riener and Kun) joined forces with a researcher from Japan and the head of the US-based research arm of the American Automobile Association (AAA) to propose another related Dagstuhl seminar. And Riener continues his collaboration with Shaer and Kun on exploring in-vehicle interfaces for work in automated vehicles – within this collaboration Kun is planning a visit to Riener's lab in the summer of 2022.

6.1.4 And one more goal – exploring best practices for remote, weeklong collaboration

This Dagstuhl seminar was conducted fully online, bringing about 30 researchers together to productively discuss the future of human-computer interaction for work in mobile environments, This was not an hour-long collaboration, or even a half-day workshop. Rather, this was a 5-day engagement with participants joining activities from all across the globe. This gave the organizers and participants an opportunity to think about best practices for organizing a key aspect of remote meetings. These meetings will likely affect our future of work deeply – work that will likely be global, collaborative and creative. Thus, we all thought about time zones (who should wake up when?), Zoom etiquette (who should talk, when, for how long?), collaboration tools (such as Miro, movie watch parties, and messaging apps), how to present information in online meetings (talks or conversations?), and how to write this report together (asynchronous vs. synchronous collaboration). We don't have the answers to all of the questions. And we know that some aspects of Dagstuhl cannot be replaced by virtual platforms (we really want to climb the hill to the castle ruins again, and we want to spend time chatting with colleagues in the music room). But this 5-day online seminar did what we have come to expect from in-person Dagstuhl seminars: it connected us to colleagues, it allowed us to come up with new ideas, and it taught us something unexpected (yes, we expect the unexpected from Dagstuhl). This unexpected things was that if we carefully organize them, we can have VERY productive, and even enjoyable, multi-day engagements online with a global group of collaborators. Thank you Dagstuhl for allowing us to learn this key lesson.

Remote Participants

- Sun Joo Ahn
University of Georgia – Athens, US
- Ignacio J. Alvarez
Intel – Hillsboro, US
- Laura Boffi
University of Ferrara, IT
- Susanne Boll
Universität Oldenburg, DE
- Stephen Brewster
University of Glasgow, GB
- Duncan Brumby
University College London, GB
- Gary Burnett
University of Nottingham, GB
- Yi-Chao Chen
Shanghai Jiao Tong University, CN
- Lewis Chuang
IfADo – Dortmund, DE
- Anna Cox
University College London, GB
- Birsen Donmez
University of Toronto, CA
- Geraldine Fitzpatrick
TU Wien, AT
- Mark D. Gross
University of Colorado – Boulder, US
- Shamsi Tamara Iqbal
Microsoft Research – Redmond, US
- Christian P. Janssen
Utrecht University, NL
- Wendy Ju
Cornell Tech – New York, US
- Andrew Kun
University of New Hampshire – Durham, US
- John D. Lee
University of Wisconsin – Madison, US
- Siân Lindley
Microsoft Research – Cambridge, GB
- Mark McGill
University of Glasgow, GB
- Helena M. Mentis
University of Maryland – Baltimore County, US
- Alexander Meschtscherjakov
Universität Salzburg, AT
- Bastian Pfleging
TU Eindhoven, NL
- Champika Manel Epa Ranasinghe
University of Twente – Enschede, NL
- Andreas Riener
TH Ingolstadt, DE
- Sayan Sarcar
University of Tsukuba – Ibaraki, JP
- Clemens Schartmüller
TH Ingolstadt, DE
- Martina Schuß
TH Ingolstadt, DE
- Orit Shaer
Wellesley College, US
- Gregory F. Welch
University of Central Florida – Orlando, US

Report from Dagstuhl Seminar 21243

Compute-First Networking

Edited by

Jon Crowcroft[1], Philip Eardley[2], Dirk Kutscher[3], and Eve M. Schooler[4]

1 **University of Cambridge, GB,** jon.crowcroft@cl.cam.ac.uk
2 **BT Applied Research – Ipswich, GB,** philip.eardley@bt.com
3 **FH Emden, DE,** dku@dkutscher.net
4 **Intel – Santa Clara, US,** eve.m.schooler@intel.com

Abstract

A Dagstuhl seminar on Compute-First Networking (CFN) was held online from June 14th to June 16th 2021. We discussed the opportunities and research challenges for a new approach to in-network computing, which aims to overcome limitations of traditional edge/in-network computing systems.

The seminar discussed relevant use cases such as privacy-preserving edge video processing, connected and automated driving, and distributed health applications leveraging federated machine learning. A discussion of research challenges included an assessment of recent and expected future developments in networking and computing platforms and the consequences for in-network computing as well as an analysis of hard problems in current edge computing architectures.

We exchanged ideas on a variety of research topics and about the results of corresponding activities in the larger fields of distributed computing and network data plane programmability. We also discussed a set of suggested PhD topics and promising future research directions in the CFN space such as split learning that is supported by in-network computing.

Seminar June 13–16, 2021 – http://www.dagstuhl.de/21243

2012 ACM Subject Classification Networks → Network architectures; Networks → Network design principles

Keywords and phrases Distributed Machine Learning, distributed systems, edge-computing, in-network computing, networking

Digital Object Identifier 10.4230/DagRep.11.5.54

Compute-First Networking, *Dagstuhl Reports*, Vol. 11, Issue 05, pp. 54–75
Editors: Jon Crowcroft, Philip Eardley, Dirk Kutscher, and Eve M. Schooler
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

1 Executive Summary

Dirk Kutscher

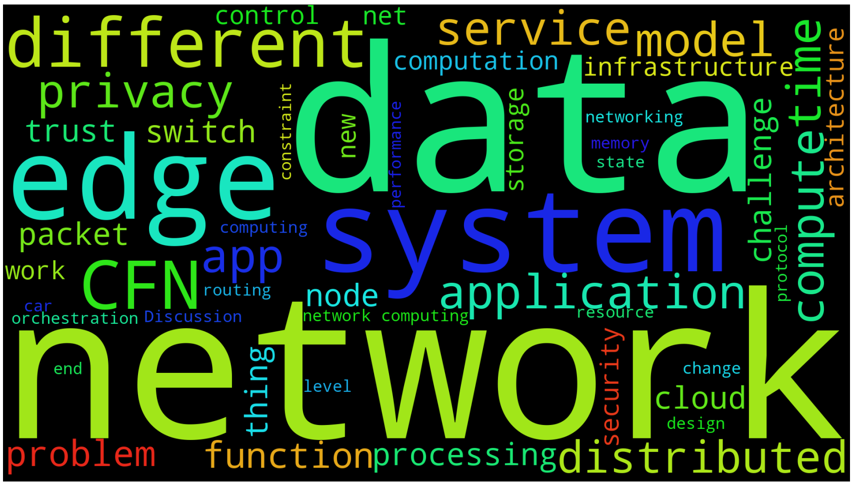

Edge- and more generally In-Network Computing are key elements in many traditional content distribution services today, typically connecting cloud-based computing to consumers. The advent of new programmable hardware platforms, research and wide deployment of distributed computing technologies for data processing, as well as new exciting use cases such as distributed Machine Learning and Metaverse-style ubiquitous computing are now inspiring research of more fine-granular and more principled approaches to distributed computing in the “Edge-To-Cloud Continuum”.

The Compute-First Networking Dagstuhl seminar has brought together researchers and practitioners in the fields of distributed computing, network programmability, Internet of Things, and data analytics to explore the potential, possible technological components, as well as open research questions in an exciting new field that will likely induce a paradigm shift for networking and its relationship with computing.

Traditional overlay-based in-network computing is typically limited to quite specific purposes, for example CDN-style edge computing. At the same time, network programmability approaches such as Software-Defined Networking and corresponding languages such as P4 are often perceived as too limited for application-level programming. Compute-First Networking (CFN) views networking and computing holistically and aims at leveraging network programmability, server- and serverless in-network computing and modern distributed computing abstraction to develop a new system’s approach for an environment where computing is not merely and add-on to existing networks, but where networking is re-imagined with a broader and ubiquitous notion of programmability.

We expect this approach to enable several benefits: it can help to unlock distributed computing from the existing silos of individual cloud and CDN platforms – a necessary condition to enable Keiichi Matsuda’s vision of Hyper-Reality and Metaverse concepts where the physical world, human users and different forms of analytics, and visual rendering services constantly engage in information exchanges, directly at the edges of the network. It can also help to provide reliable, scalable, privacy-preserving and universally available platforms for Distributed Machine Learning applications that will play a key role in future large-scale data collection and analytics.

CFN's integrated approach allows for several optimizations, for example a more informed and more adaptive resource optimization that can take into account dynamically changing network conditions, availability of utilization of compute platforms as well as application requirements and adaptation boundaries, thus enabling more responsive and better-performing applications.

Several interesting research challenges have been identified that should be addressed in order to realize the CFN vision: How should the different levels of programmability in todays system be integrated into a consistent approach? How would programming and communication abstractions look like? How do orchestration systems need to evolve in order to be usable in these potentially large scale scenarios? How can be guarantee security and privacy properties of a distributed computing slice without having to rely on just location attributes? How would the special requirements and properties of relevant applications such as Distributed Machine best be mapped to CFN – or should distributed data processing for federated or split Machine Learning play a more prominent role in designing CFN abstractions?

This seminar was an important first step in identifying the potential and a first set of interesting new research challenges for re-imaging distributed computing through CFN – an exciting new topic for networking and distributed computing research.

2 Table of Contents

3 Overview of Talks

3.1 In-network telemetry

Gianni Antichi (Queen Mary University of London, GB)

Abstract

The possibility to easily add new functionality to network data planes has lately opened new exciting research directions towards understanding how such a programmability can impact the design of networks as well as their services. This talk focussed on network monitoring and analyzed how both flow-based and packet-level telemetry can be efficiently implemented on state-of-the-art programmable switches. Starting from the limitations of current sketch-based solutions and in-band network telemetry (INT), the presentation discussed new ways to provide faster flow level analysis to network collectors and lighter in-band telemetry through the use of network coding techniques. The talk concluded with considerations on what are the next steps in the field.

Discussion Summary

In-network telemetry was discussed as an example for a dataplane mechanism that could enable several useful measurement and management services in CFN networks such as exposing dynamic load/utilization information for joint optimization and execution profiling. Depending on the use case there may be different requirements with respect to the fidelity and completeness of the observed information. The current in network telemetry concepts are interesting due to their "end-to-end" nature, not requiring state on network elements. Several references to related concepts were made including [1], [2], [3], [4], [5].

References

1 CLOUD COMPUTING TECHNOLOGY LAW AND POLICY, Microsoft. `http://www.mccrc.org/`
2 Identifying and Supporting Financially Vulnerable Consumers in a Privacy-Preserving Manner: A Use Case Using Decentralised Identifiers and Verifiable Credentials, Tasos Spiliotopoulos et al, 2021. `https://arxiv.org/abs/2106.06053`
3 Securing software by enforcing data-flow integrity, Miguel Castro et al, OSDI 2006. `https://www.microsoft.com/en-us/research/publication/securing-software-by-enforcing-data-flow-integrity/`
4 NSA, Andreas Eschbach, 2020. `https://www.luebbe.de/luebbe-belletristik/buecher/sonstige-belletristik/nsa-nationales-sicherheits-amt/id_7539196`
5 Split Learning: Distributed deep learning without sharing raw data, MIT project page. `https://www.media.mit.edu/projects/distributed-learning-and-collaborative-\learning-1/overview/`

3.2 CFN for Health

Jon Crowcroft (University of Cambridge, GB)

Abstract

Sensing well-being and health via mobile devices is becoming mainstream. Logging the data centrally doesn't scale, and causes privacy (and ethical and legal) problems. Local processing, and in-network aggregation can support the individual (life logging, diagnostics etc) and public health (e.g. epidemic, e.g. avoiding infection, and environmental risk, e.g. air quality). As a use case for CFN, and an instance of federated machine learning and a challenge for privacy designs, the system architecture for such apps make an informative study.

Discussion Summary

We discussed how covid apps for contact tracing use a compute-first networking approach. In general, aggregated information is good enough for many applications. Linking back to the previous talk, we discussed the need to be able to explain the privacy and security properties of a system, and how it is easy to imagine dystopian futures.

3.3 New Network for Distributed Systems

Jianfei He (City University – Hong Kong, HK)

Abstract

An infrastructure like Networks should be designed with stable functions which are decoupled from the semantics of specific applications. This talk proposes to address the common requirements of various distributed systems, including collective communication and consistency. It is also emphasised that new network diagrams should be justified not only by its improvement on end-to-end performance, but also by its superiority over potential solutions at upper layers.

Discussion Summary

This talk led to a discussion of generality and modularity, i.e., which are the right application-agnostic building blocks? Would it make sense to have a taxonomy of different distributed application types and their requirements for in-network and distributed computing support functions? Whereas some fundamentals of distributed computing are well-understood and will be relevant in the future, it is not always useful to predict how certain mechanisms will be used and what optimization techniques might emerge, which seems to suggest a rather flexible approach to layering and re-use. In this context, Scaling Replicated State Machines with Compartmentalization [1] and CoFlow (`https://amplab.cs.berkeley.edu/tag/coflow/`) were mentioned.

References

1 Secure Scuttlebutt, project pages. `https://scuttlebutt.nz/`

3.4 Networking matters for storage systems

Micchio Honda (University of Edinburgh, GB)

Abstract

Although networks have been programmable and able to perform computation, the trend of end-to-end encryption and stream-oriented TCP protocol make it difficult to benefit from such programmable networks. This talk presents Prism as an opportunity of significantly improving cluster resource utilization and reliability with the aid of programmable networks although the communication occurs over regular TCP and TLS. We instantiate Prism as an scale-out object storage system that communicates with the client using widely used S3 protocol.

Discussion Summary

We discussed comparison to other approaches which use custom UDP-based network protocols in terms of achieving compute-first networking and usability. We also discussed other system-level approaches such as VM live migration.

3.5 Privacy Preservation in Edge Video Processing

Jag Minhas (Sensing Feeling – London, GB)

Abstract

In today's world of virtual computers, high speed networking and algorithms powered by deep neural networks, application developers have a wide choice of architectures and arrangements at their disposal to implement working solutions. For those implementing vision processing systems in the area of human surveillance and behaviour these architectural choices will often have an impact on how end-user rights to privacy are upheld. This talk summarises an established "privacy affirming" framework for distributed vision processing systems and discusses the implications of this approach as our architectural choices widen further in a "Compute First Networking" future.

Discussion Summary

Privacy is agreed as a critical topic for compute-first networking. Distribution of processing offers the potential for improved privacy protection, for example by distributing machine learning inference to multiple nodes, but there are many subtle issues (for example, the assumption that you can trust that the system is doing what it's supposed to). There can be privacy implications even with no (other) human in the loop, for example the use of captured video for ad tracking. A full solution needs to consider the legal and policy framework as well as the technical design.

3.6 Compute Centric Networking for Connected and Automated Driving

Naresh Nayak (Robert Bosch GmbH – Stuttgart, DE)

Abstract

In-vehicular networks are rapidly evolving to support the future requirements of the automotive industry stemming, e.g., from highly autonomous driving. The introduction of microprocessor platforms (aka vehicle computers) networked over high bandwidth Ethernet backbone coupled with wireless connectivity has transformed the modern day vehicle into an IoT device on wheels. The continually increasing amount of computing capability within a vehicle, however, comes at a cost (monetary, energy, etc.). In-network computing is a promising avenue to include the computing capabilities from the infrastructure, and thus, reduce the corresponding costs. This talk presented different use cases where driving functions could be offloaded from a vehicle to the infrastructure and discussed the associated research challenges.

Discussion Summary

A vehicle now contains many IoT devices and produces a lot of data (raw data at 10s TB/hour). So there can be a compute-first network within the car, as well as in the nearby roadside infrastructure. Some of the use cases have critical timing requirements (for safety reasons) – so may be a use case for deterministic computing – whilst others are more tolerant of latency. Security and trust are critical issues, especially if cars use information from other cars as part of automated driving.

3.7 EVEn Harder Challenges

Erik Nordmark (Zededa – San Jose, US)

3.7.0.1 Abstract

The LF Edge EVE project targets IoT workloads at the smart device edge whilst LF Edge has a suite of other projects for other parts of the edge-cloud continuum. EVE's approach to control issues is that there is a single controller per edge node with a "phone home" function. EVE's approach and assumptions to issues such as trust, placement and security were outlined.

3.7.0.2 Discussion Summary

Discussion was about detailed issues concerning the LF Edge EVE project.

3.8 Networking and Computing – the great confluence

David Oran (Network Systems Research & Design – Cambridge, US)

Abstract

Two contemporaneous trends are merging the two separate yet interdependent technologies of computing and networking. Hardware and software historically associated with computing complexes (General-purpose CPUs running virtual machines, conventional operating systems, and languages) are being used more and more to host networking functions at all but the highest speed tiers. Networking devices such as switches, routers and NICs are becoming programmable in ways that allow general purpose computing to be done "in the network". This talk examines these trends, presents some of the salient research illuminating the advantages and limitations, and speculates on where this merging of technologies might take us.

Discussion Summary

Comparing the capabilities of servers and switches, for example in terms of cycles/bit, memory and rate of feature evolution: server capabilities are superior (providing only a limited number of ports are needed) and programming is easier.

Many topics were discussed.

One issue associated with joint optimisations (e.g, considering networking, compute, storage and possibly energy resources) is if the coordinated allocation of resources is allowed across boundaries. We should learn from the experiences with multi-provider QoS which has been a non-starter for decades (for commercial and policy reasons – we do have the necessary protocols).

3.9 Split Learning Opportunities with CFN

David Oran (Network Systems Research & Design – Cambridge, US)

Basics of Split Learning Summary of potential advantages for combining split learning with edge computing 3 Possible use cases, focused on IoT (mostly)

3.10 (Some Random Thoughts on) In-Network Compute Architecture

Jörg Ott (TU München, DE)

Abstract

In-network computing, the execution of program logic on behalf of an application "inside" the network, may be interpreted in many ways, yielding rich history: from the early work on active networking to flavors of mobile offloading in the context of pervasive computing to

programmable data planes to edge computing to recent extensions to information-centric networking in support of computations, among others. These interpretations differ in the complexity of the application structure exposed to "the network", the code complexity, where and on what kind of devices program logic is executed, and how and when the location(s) for execution are determined and routed to. In this talk, we explore some architectural considerations on in-network computing. We start out from basic networking/computing primitives and explore basic units of computation and interaction, look into different application models, where the "main" control thread (if any) could reside, and how traditional Internet design principles, especially best effort operation and the end-to-end principle, could translate to in-network computing. Another important question is how much abstraction from (as opposed to insight into) the execution scheduling and orchestration a network should provide. We discuss the importance of naming and scoping concepts for in-network functions (and state) that make up the program logic: for invoking them from applications outside the program logic as well as for the possible interactions among multiple (distributing) functions comprising this very logic. We look at name resolution and routing mechanisms and how transport layer mechanisms relate to in-network computing. Overall, there seem to be many feasible approaches, and many design decisions likely depend on how such in-network compute infrastructure would eventually be realized: augmenting IP, running as overlays, or defining a novel clean-slate network architecture.

Discussion Summary

The presentation raised many questions that prompted much discussion and further questions, including: whether one needs to think about specific use case examples and the business models; what the API should be like; what do we name (nodes, data, functions, processes etc) – do we explicitly invoke a function, or do we address an object and the function happens consequently; "ownership" of a function's execution and the associated resource; partial updating is normal; how to take account of the unreliability of memory and so also how to degrade gracefully; and whether all computation should be explicit (so it is all "end to end", with no implicit or "hidden" computation).

3.11 Deconstructing and Reimagining Orchestration

Andy Reid (BT – Ipswich, GB)

Abstract

Orchestrators are a general concept, including SDN controllers, cloud controllers and package managers, which are layered (one form of orchestrator feeds another). Each has an interface through which its user can inject policy. Each includes the same set of activities (such as parser, translator, configuration manager and optimiser). Intent-based networks have the same general approach. Today's orchestrators are focussed on the management plane. The proposal is that it would be beneficial to re-style orchestrators as primarily part of the control plane, which would imply a greater degree of automated operation, with security built into the protocol. One implication suggested is that functions should be called by name, not by address.

Discussion Summary

One topic of discussion was about optimisation: whether this is easier with hard or soft constraints, given the overall system is likely to be non-linear and a "soft" answer is wanted. Since compute resource has no real equivalent of the "bit" as a unit of networking resource, it is harder to optimise. Another topic was how to evolve from the current state to the desired state (as recommended by the optimiser), without putting the network /system through an unsafe intermediate state (to avoid this problem, routing policy changes have to go through many intermediate steps). A further evolution issue is that feedback is needed about whether the resources required in the desired state are actually available (especially as there may be inaccuracies), whereas today's control planes typically just "try now" – the potential role of digital twins was mentioned.

3.12 Heterogenous Networks and Laying out the Compute Graph

Peer Stritzinger (Peer Stritzinger GmbH – Maisach, DE)

Abstract

For use cases in edge computing and IoT placing compute functionality in a heterogeneous not fully connected network is an open challenge. Focussing on a computational model consisting of lightweight processes and message passing approaches to modelling distributed algorithms and compute and network resources have been explored. Process and messages are a model to which all distributed computational models can be mapped easily. Relatively static vs. highly dynamic mapping is necessary for different use cases. Properties of processes vs. nodes and messages vs. links need to be combined with topology of communication and network. Orchestration needs to be implementable in a distributed manner without central control.

Discussion Summary

Discussion included Erlang approaches to achieving this distributed computation model. Other issues raised included: the potential role of time sensitive networking and deterministic networking; energy and green energy sensitivity; 'interference' from other traffic or computations; and how to scale to distributed IoT systems that could have billions of processes.

3.13 After the Fun the Hard Stuff

Christian Tschudin (Universität Basel, CH)

Abstract

A core assumption of traditional distributed computing is that processes send messages among themselves to coordinate their work while in this talk we suggest exploring memory-mediated coordination. The Internet is process-centric (packet switching among protocol entities)

which is typically made available to applications in the form of interprocess communication (IPC). Decentralized systems like Secure Scuttlebutt, however, have no notion of IPC and instead rely on any form of content propagation (id-centric append-only logs). Such a content-centric venue may become important when we envisage highly volatile in-network computing environments where processes are not stable enough to receive and answer messages, or because they follow probabilistic computing models, or they simply are not concurrently online. We exclude approaches based on a request-reply pattern à la CCN/NDN and Named-Function Networking (which was fun, but too perfect) because they imply some notion of processes. While one could try to harden the network fabric to provide message buffering beyond process lifetime (mailbox systems for actors), we suggest to focus on bare in-network memory and force processes to coordinate via "DIstributed Memory Side Effects" (DIMSE) only. Such an architecture would have immediate benefits beyond crash-resilience, namely for delay tolerance, transport agnosticity and even for data-center networks (fast remote memory access for HPC).

Discussion Summary

The talk prompted a very lively discussion. The idea is for push-only data distribution, with social forwarding-and-filter (meaning that filtering is based on the social graph). The role of "scent marks" to enable this was discussed, with each node in memory storing a collection of scent marks. Questions were raised about garbage collection, contracts and motivations (and tit-for-tat issues), and lessons from earlier systems such as Linda, SRM and floating content.

3.14 In-network computing challenges and opportunities

Noa Zilberman (University of Oxford, GB)

Abstract

Programmable network devices have changed the way we perceive the network, allowing us to rethink where, when and how we process data. The emergence of in-network computing, the execution of standard host applications within network devices, has been especially promising. However, the rise of in network computing has brought with it multiple challenges, ranging all the way from architecture challenges, through micro-architecture, design and validation, to multi-tenancy. This talk explored the many challenges in designing functional, high performance, and robust in-network applications. These challenges propose new research opportunities, allowing to turn in-network applications from research-ideas into commercially-viable solutions.

Discussion Summary

The discussion focused on the applicability of P4 to distributed computing. Whereas P4 dataplane programming is sometimes considered as the most promising in-network computing approach, this talk highlighted many of the practical difficulties. Whereas some point solutions for distributed computing applications and optimization have been developed, such as [1], the lack of Turing completeness and the vast heterogeneity in platform capabilities seems to suggest that P4 dataplane programming would be suitable for certain packet-level optimization while possibly integrated into a larger, heterogeneous distributed computing framework. In that sense, P4 dataplane programming could represent a certain class of enabling platforms, and a compelling research challenge might be to productively leverage such special environments. Similar concerns could arise for other powerful, but really specialized environments such as Multi-GPU communication (e.g., https://www.nvidia.com/en-gb/data-center/nvlink/).

References

1 Raft does not Guarantee Liveness in the face of Network Faults, Heidi Howard et al, 2020. https://decentralizedthoughts.github.io/2020-12-12-raft-liveness-full-omission/

4 Open problems

4.1 Potential Research Problems

All Participants

The seminar held one session on discussing research opportunities and potential PhD topics in the CFN (Compute First Networking) field to develop an understanding of what original and substantial research activities could look like. Here is a summary of the ideas discussed:

- Dynamic compilation and placement of CFN instances – network-wide, cross-platform, just-in-time compilation
- Profile-Guided Optimization as CFN Optimizer – Profile-guided optimization is an existing technique that uses profiling test runs of a program to identify parts of the code that are executed more frequently and should be optimised. The idea would be to use data plane telemetry to do the profiling.
- Layer 3 Protocol for In-Network Computing – a clean slate protocol is suggested for routing to compute, possibly building on some of the ideas from ICN and NFN
- In-Network Computing for Network Management – with a logically centralised control plane but policy distributed to each node
- Analysis of CFN – the research topic would be to demonstrate the "provable" gain for CFN
- In-Network Edge Computing for Smart Transportation – the goal would be fast, safe and improved coordination of all sorts of transport users (cars, lorries, trams, cycles, pedestrians, buses, and so on). Questions include the role of road side units and direct coordination between vehicles, and privacy-preserving techniques and data ownership.
- Climate change – what role can CFN play in the goal of net zero?
- Faster packet processing after hitting specialization limits – packet processing speeds are now improving only slowly, and specialisation for a given domain works once, so what are the options to improve network packet processing speed? The potential roles of reconfigurable hardware, specialist hardware, specialist channel, and quantum computing.
- Distribution Usage Descriptions to deny denial of service – the idea would be to infer about flows and to police them in a decentralised manner.
- Gracefully Degradable Turing Machines – faster hardware is hitting reliability limits, so algorithms are needed that cope with memory faults

- New Network for Distributed Systems – routers as we know them will decrease / disappear, instead we will have compute and inter-computing – in that world, what are the basic functions?
- Planarian Distributed Computing – the notion of a 'breaking point' in distributed computing – partitioning and merging. Is there an interesting space between consensus based and eventual consistency partition tolerance?
- Towards a Formalism for CFN – Start with Misra and Chandy's UNITY formalism, which has distributed computations as collections of iterated assignment statements, and extend into fault tolerance and different models for the comms/compute/storage
- Split Learning Opportunities with CFN – apply Deep Neural Networks, but split vertically across layers rather than horizontally across samples, and smashing the output data at the boundary to preserve only relevant excitation coefficients. This can help with privacy, by confining data to the device only, or in secure enclaves in edge computing nodes; it also can provide failure resiliency, as there's degraded operation (rather than none) when the cloud back-end can't be reached.

5 Findings and Recommendations for Compute First Networking

All Participants

In this section we present the organisers' conclusions from the seminar. Undoubtedly not all participants will agree with all of them.

The field of Compute First Networking (CFN) is at an exciting moment. We speculate that compute-first networking is at the same stage as cloud computing 20 years ago. At that moment, we had some insight about both the commercial drivers (the use cases for utility computing) and the core technical ideas (data centres, virtualisation and fairly fast networks). But we didn't understand how so many industries would be transformed, nor all the technical work required in order to be able to deploy and orchestrate services (such as map-reduce, stream processing, load balancers, migration, containers, and distributed ledgers).

We believe that Compute First Networking has similar potential to transform the business world and the way people live. Today, we have the first ideas and partial understandings about the use cases, commercial drivers and technical directions, but there is a huge amount still to do.

In the absence of new service models, it seems likely that the hyperscalers will extend their dominance of cloud computing into CFN. For those who would like to see a more diverse ecosystem, now is therefore an opportune moment to invest in researching the key components needed to disrupt the status quo. We recommend that government and industry fund extensive research and development in order to help bring CFN to fruition, and allow fresh innovation by a decentralised gallery of players.

5.1 Use Cases and Business Drivers for CFN

The continued consolidation in data centre clouds has made other deployment models appear as an exception – for example, we talk about edge computing when referring to non-cloud-only applications (where most of the artefacts and compute resources are concentrated in the cloud).

Non-cloud-only applications are becoming increasingly important, due to mainstream consumer-focused applications (scalability of content distribution and low-latency requirements for interactive applications) but specifically due to new applications that leverage the availability of Internet-connected physical resources such as sensors and video cameras, and the potential advantages of integrating these into larger distributed applications.

For example, from a mere data logistics perspective, (semi)-autonomous vehicles produce large amounts of data, mostly destined for local consumption (video and various sensors), but increasingly also for external processing and storage – for tele-driving, diagnostics, forensics and other applications. Such applications fundamentally require a robust data offloading and processing infrastructure – to some extent the reverse functionality of today's downstream CDNs (Content Distribution Networks). Processing should enable filtering, aggregation, and other forms of transformation, and must thus be fully programmable.

Not all interactions are upstream-data-based – for example, car and factory networks have rigid reliability and low-latency requirements. Corresponding control loops may be driven by local data as well as by remote information and instructions. Enabling such robust control in the presence of dynamic and heterogeneous sub-system properties calls for infrastructure support that enables operating distributed applications across different failure and time domains .

Infrastructure programmability and resource multiplexing raise the question of multi-tenancy, i.e., the ability to run several workloads over shared computing and network infrastructure with adequate isolation . Fundamental building blocks such as suitable encryption algorithms, trusted execution, attestation etc. exist and are emerging. They fuel the design of more reliable and trustworthy distributed applications over shared infrastructure, enabling new use cases.

When edge sensor data contains personally identifiable information (PII) or can be used to construct PII by combining multiple inputs, additional concerns regarding privacy preservation arise. We reject the suggestion that privacy is over-rated – there are already too many dystopian abuse scenarios – and we reject the 'counsel of despair' that it is already too late (because the hyperscalers and some governments, service providers, and/or credit card companies collect so much information today). On the other hand, the Covid-19 pandemic has shown importance of finding a 'middle way', where health data can be leveraged for tracking, trend analysis, and understanding transmission patterns. We believe that a future system should preserve privacy by design , meaning that it technically prevents leakage of PII (such as raw video data). This could enable a whole new class of applications, such as the use of medical data in a productive and ethical manner. In-network computing actually offers some promise for new approaches to privacy preservation, because computation is intrinsically distributed, and end users might be able to control which functions/analytics are run, where they run (in the local network), for whom, and on what data.

The economics of compute first networking merits research work. What types of compute function are likely to be of most value? Are there new use cases where compute-first networking offers particular benefit? What applications are distributed and multi-party? What are the externality costs like power, real estate, right-of-way, regulation of physical infrastructure?

Also worthwhile is consideration of the legal aspects of compute first networking. At a fundamental level, should there be regulation to achieve "in-network compute neutrality"? Today, some countries have regulations about "net neutrality" for the networking, but not about the cloud compute, which has (probably) been key to enable the development of a plethora of applications, but has disintermediated the telecoms operators from the value chain. What is the preferred societal outcome for CFN? There are also detailed questions to be resolved, such as the appropriate legal frameworks when it may be necessary to reveal some of the data or its analysis. Examples include accident data for cars, medical malpractice, and criminal conspiracy.

5.2 Technical challenges for CFN

While distributed computing functionality in the network is an essential design feature in many mainstream applications today, current architectural models can only provide a limited form of in-network distributed computing due to a lack of appropriate programming models and run-time visibility.

Many applications today are becoming increasingly multi-party and distributed, e.g., user front-end applications talking to a microservice backend, possibly mediated by proxies, load balancers etc. However, from an application developer perspective, it is difficult and possibly undesirable to make the network fully transparent to the programmer: there are performance inhomogeneities in both throughput and delay, complex partial failures that require some form of response, depending on the distributed system's run-time behaviour. With current systems, except for rather static directed acyclic graph structures (such as Dataflow systems), traditional programming models today only easily exploit localized parallelism, so it is still not possible to reason about wider system parallelization without requiring a lot of orchestrator intelligence.

DevOps-based systems can theoretically enable incremental partial deployment and dynamic module updates, however coordination with network underlays is not trivial and not always possible due to strict separation into organizational and protocol realms.

In addition, computing and communications are on different cost/performance trajectories , which results in additional platform heterogeneity and different typical functions that one would run on compute servers vs. network switches. From a processing and memory perspective, compute servers can afford to spend many cycles per processed bit and to perform memory intensive operations, whereas network switches are concerned with line-rate processing on optimized hardware, allowing only a few cycles per bit with only small or moderate memory consumption. Correspondingly, the workloads are also quite different, especially considering the need for multi-tenancy and tenant isolation (requiring cryptographically enabled isolation on servers, which is typically not available on network switches). Thus, we believe that there will be "puddles and ponds" of compute, connected naturally via networks (rather than every device including networking and compute). For scalability reasons these resources are more likely to congregate.

Programmable data planes are likely to be an important component, but today's technologies, such as P4, are not fully formed as yet, as their capabilities are too limited and there are many tricky issues remaining to be solved for generic in-network computation. In addition to generic ('universal') processors, there will be a need for heterogeneous hardware specialised to jobs such as sensor processing and filtering (e.g. sparse FFT), or query minimization for differential privacy.

On privacy , a first (but not sufficient) technical step could be the careful selection of processing loci in the system, i.e., perform data analytics close to the data source and then only share processed, anonymized result data. Federated machine learning could be one such pattern that enables the controlled sharing of the results of analytics, for example featuring individual health data logging and public health data sharing via user-controlled analytics functions. Other approaches, such as split learning could map the different deep learning layers onto different nodes/domains in the network, thus enabling a privacy-preserving distributed training process, where each layer can only look at its own raw data.

Provenance tracking is also important, as insights derived from data are dependent on data trustworthiness. Also, on a compute node, isolation is needed between tenants, with some kind of trusted computing.

A key research topic is about the placement of functions, distributed across the network (the processing graph), and the related optimisation of usage of resources . This requires joint consideration of, and trading off between, different types of resource capabilities – compute, storage, networking, (green) energy – as well as other constraints, such as time-sensitivity of the application, and the users' appetite for risk and cost. To take a simple example: it may be better to move the function to be near the data, rather than shift all the data to a function running in the central cloud or an edge cloud, particularly if the volume of data produced makes it impossible to move in its original form. Since the optimisation is an NP-complete problem, it is not about finding a unique global optimum, but a reasonably good enough solution. The optimisation also has to take account of some additional challenges. One is how to include resilience as a key criteria (since producing densely connected edge networks is probably prohibitively expensive) — including solutions that are stable for some degree of failures, and the ability to safely transition from one solution to another without violating service level criteria, or via a risky, unstable intermediate point.

Another challenge is how to include the cost of change — including the time to compute the solution vs. the time before the configuration has to change (due to changes in the workload and/or the resource availability), the cost of spinning resources up and down, and – on a longer timescale – the optimisation of value (rather than cost) through time. A third issue is the inclusion of policy preferences and constraints of the CFN operator and the users. Finally, it is possible that the resources may be contributed by different owners (for example, owners of the RAN, access network, distributed compute, compute in a POP), whilst spot price manipulation should be prevented.

Another topic fruitful for research could be to re-consider the balance of the management and control planes . Traditionally the control plane operates automatically and autonomously, driven by protocols and algorithms, whilst the management plane operates with manual intervention and specialisation. Can orchestration, in its broadest sense, be re-architected so that it is a control plane set of functionality that automatically converts the "intent" of the users (and networks) into components instantiated on the hosting infrastructure?

Finally it is worthwhile to re-consider earlier approaches that were abandoned, in light of compute-first networking. For example, I3, active networking, and network coding.

In summary, today the building blocks and concepts for CFN are only partially understood. There is great potential and significant challenges to achieve CFN as a platform that can enable new use cases and business opportunities.

References

1 An Overview of Privacy in Machine Learning Emiliano De Cristofaro, March 2020. https://emilianodc.com/PAPERS/privacyML.pdf
2 Rearchitecting Kubernetes for the Edge, Andrew Jeffery et al, EdgeSys 2021. https://arxiv.org/abs/2104.02423
3 The ephemerizer: Making data disappear, Radia Perlman, 2005. https://www.researchgate.net/publication/228360589_The_ephemerizer_Making_data_disappear

4 Chaffing and winnowing, Wikipedia.
https://en.wikipedia.org/wiki/Chaffing_and_winnowing
5 How mass testing helped limit the spread of COVID-19 at the University of Cambridge, Craig Brierley, 2020.
https://www.cam.ac.uk/stories/screeningprogramme
6 On the Utilisation of Persistent Programming Environments, Richard Cooper, PhD thesis, 1989. http://theses.gla.ac.uk/77963/
7 Membrane computing, Wikipedia. https://en.wikipedia.org/wiki/Membrane_computing
8 Opportunistic content sharing applications, Ott et al, NoM12. http://www.netlab.tkk.fi/~jo/papers/2012-nom-floating-api.pdf
9 Efficient Publish/Subscribe-Based Multicast for Opportunistic Networking with Self-Organized Resource Utilization, Greifenberg et al, AINAW08. https://dl.acm.org/doi/10.1109/WAINA.2008.255
10 A reliable multicast framework for light-weight sessions and application level framing, Floyd et al, ACM CCR 1995. https://dl.acm.org/doi/10.1145/217391.217470
11 Manufacturer Usage Description Specification, IETF, 2019. https://datatracker.ietf.org/doc/html/rfc8520
12 Automated valet parking video:https://www.youtube.com/watch?v=0q-kLX5ISZA
13 An Advanced Teleoperation Testbed, oss et al, 2008. https://link.springer.com/chapter/10.1007/978-3-540-75404-6_28
14 Autocloud, or how to use your car battery to save energy! Jon Crowcroft. https://www.cl.cam.ac.uk/~jac22/talks/autocloud.pdf
15 How Internet Concepts and Technologies Can Help Green and Smarten the Electrical Grid, Keshav et al, 2010. http://svr-sk818-web.cl.cam.ac.uk/keshav/papers/10/greennet.pdf
16 Modelling Incentives for Collaboration in Mobile Ad Hoc Networks, Crowcroft et al, WiOpt'03. https://hal.archives-ouvertes.fr/inria-00466747/
17 In-situ OAM IPv6 Options, Bhandari et al, work in progress. https://datatracker.ietf.org/doc/html/draft-ietf-ippm-ioam-ipv6-options
18 IOAM (IPv6) in Linux kernel, Iurman. https://datatracker.ietf.org/meeting/105/materials/slides-105-ippm-5-ioam-implementation-00
19 Scalable Network Management Using Lightweight Programmable Network Services, Calvert, 2006. https://link.springer.com/article/10.1007/s10922-005-9013-6
20 Pervasive Debugging, Hand et al. https://www.cl.cam.ac.uk/research/srg/netos/projects/archive/pdb/papers/2004.escience.pdf
21 Generic Multicast Transport Services: Router Support for Multicast Applications, Cain et al. 2000. https://www.semanticscholar.org/paper/Generic-Multicast-Transport-Services%3A-Router-for-Cain-Towsley/aecc9f22d0b0dfd7b416a1837d051bfb92a4a27a
22 Tiny Packet Programs for low-latency network control and Monitoring, Jeyakumar et al, Hotnets 13. https://conferences.sigcomm.org/hotnets/2013/papers/hotnets-final47.pdf
23 Probabilistic Synchronous Parallel, Wang et al, 2017. https://arxiv.org/abs/1709.07772
24 NaaS: Application-Specific Network Services, project page. https://lsds.doc.ic.ac.uk/projects/naas
25 The Case For In-Network Computing On Demand, Tokusashi et al, Eurosys 19. https://dl.acm.org/doi/10.1145/3302424.3303979
26 P4xos: Consensus as a Network Service, Dang et al, ToN 2020. https://ieeexplore.ieee.org/abstract/document/9095258

27 Partitioned Paxos via the Network Data Plane, Dang et al, 2019. https://arxiv.org/abs/1901.08806
28 From photons to big-data applications: terminating terabits, Zilberman et al, 2016. https://royalsocietypublishing.org/doi/full/10.1098/rsta.2014.0445
29 Taurus: An Intelligent Data Plane, Swamy et al, 2020. https://arxiv.org/abs/2002.08987
30 CHANGING CONDITIONS FOR NEURAL NETWORK PROCESSING, Hemsoth, 2020. https://www.nextplatform.com/2020/04/07/changing-conditions-for-neural-\network-processing/
31 IPU processors, webpaeg. https://www.graphcore.ai/products/ipu
32 NVLink and NVSwitch, webpage. https://www.nvidia.com/en-gb/data-center/nvlink/
33 Introduction to Actors, webpage. https://doc.akka.io/docs/akka/current/typed/actors.html
34 Fibbing: Central Control over Distributed Routing, Tilmans et al, 2016. https://datatracker.ietf.org/meeting/97/materials/slides-97-rtgarea-\fibbing-anrp-00
35 What is Autonomic Computing? Ghosh, 2020. https://thecustomizewindows.com/2020/04/what-is-autonomic-computing/
36 Rate of convergence of increasing path-vector routing protocols, Daggitt et al. https://www.cl.cam.ac.uk/~tgg22/publications/icnp_2018.pdf
37 Prism: Proxies without the Pain, Hayakawa et al, NSDI 21. https://www.usenix.org/conference/nsdi21/presentation/hayakawa
38 PASTE: A Networking Programming INterface for Non-Volatile Main Memory, Honda et al, NDSI 18. https://www.usenix.org/conference/nsdi18/presentation/honda
39 Bigraphs and Their Algebra, Milner, 2008. https://core.ac.uk/download/pdf/82274932.pdf
40 Modelling and verification of large-scale sensor network infrastructures, Sevegnani et al, ICECCS 2018 https://ieeexplore.ieee.org/document/8595061
41 Compute First Networking: Distributed Computing meets ICN, Krol et al, ICN 2019. https://dl.acm.org/doi/10.1145/3357150.3357395
42 XenoSearch: Distributed Resource Discovery in the XenoServer Open Platform, Spence et al. https://www.cl.cam.ac.uk/research/srg/netos/papers/2003-xenosearch.pdf
43 Resolution strategies for serverless computing in information centric networking, Scherb, PhD thesis 2020. https://edoc.unibas.ch/77823/
44 AntidoteDB, website. https://www.antidotedb.eu
45 MobilityFirst Future Internet Architecture Project, Seskar et al, 2011. https://www.researchgate.net/publication/220785899_MobilityFirst_Future_Internet_Architecture_Project
46 Collective communication, Message Passing Interface Forum, 1995. https://www.mpi-forum.org/docs/mpi-1.1/mpi-11-html/node64.html#Node64
47 Coflow project website. https://amplab.cs.berkeley.edu/tag/coflow/
48 Scaling Replicated State Machines with Compartmentalization, Whittaker et al, 2020. https://arxiv.org/abs/2012.15762
49 Marc Brooker's blog. https://brooker.co.za/blog/
50 Amazon S3 Object Lambda website. https://aws.amazon.com/s3/features/object-lambda/
51 Development and management of collective network and cloud computing infrastructures, Viñas, PhD 2019. https://dsg.ac.upc.edu/rogerb-phd

52 Age of Information: An Introduction and Survey, Yates et al, JSAC 2021. https://ieeexplore.ieee.org/document/9380899
53 The PGM Reliable Multicast Protocol, Gemmell et al. https://www.cl.cam.ac.uk/teaching/1213/R02/papers/pgm_ieee_network.pdf
54 MLIR: A Compiler Infrastructure for the End of Moore's Law, Lattner et al, 2020. https://arxiv.org/abs/2002.11054
55 In-Network Computing for App-Centric Micro-Services, Trossen et al, 2021. Work in progress. https://datatracker.ietf.org/doc/html/draft-sarathchandra-coin-appcentres
56 Milking the Cache Cow With Fairness in Mind, Wang et al, IEEE Transactions on Networking, 2017. https://ieeexplore.ieee.org/document/7945473
57 What the Fitbit Saw During That 4.4 Quake, Miller, KQED 2018. https://www.kqed.org/futureofyou/438270/what-your-fitbit-and-smartphone-saw\-during-that-4-5-quake
58 On the duality of resilience and privacy, Crowcroft, Proceedings of Royal Society, 2014. https://royalsocietypublishing.org/doi/pdf/10.1098/rspa.2014.0862
59 Critical regions and region-disjoint paths in a network, Trajanovski et al, IFIP 2013. https://ieeexplore.ieee.org/document/6663507
60 Climate change: science and solutions, Royal Society, 2021. https://royalsociety.org/topics-policy/projects/climate-change-science-solutions/
61 Programmable Network Data Planes, Antichi et al, Report from Dagstuhl Seminar 19141, 2019. https://drops.dagstuhl.de/opus/volltexte/2019/11295/pdf/dagrep_v009_i003_p178_19141.pdf
62 A Foundation of Parallel Programming, Misra, 1988. https://www.cs.utexas.edu/users/misra/psp.dir/Marktoberdorf-88.pdf
63 On chemical and self-healing networking protocols, Meyer. PhD thesis, 2011. http://edoc.unibas.ch/diss/DissB_9383
64 Perspective: Energy Landscapes for Machine Learning, Ballard et al, 2017. https://arxiv.org/pdf/1703.07915.pdf
65 ACM FAccT conference (Fairness, Accountability and Transparency) https://facctconference.org/index.html
66 Counterfactual Fairness, Kusner et al, 2018 https://arxiv.org/abs/1703.06856
67 Path-Specific Counterfactual Fairness, Chiappa et al, 2018. https://deepmind.com/research/publications/path-specific-counterfactual-fairness
68 Liquid Data Networking, Byers et al, ACN 2021 https://dl.acm.org/doi/10.1145/3405656.3418710
69 Jayadev Misra; A Foundation of Parallel Programming; https://www.cs.utexas.edu/users/misra/psp.dir/Marktoberdorf-88.pdf
70 Vimalkumar Jeyakumar, Mohammad Alizadeh, Changhoon Kim, and David Mazières. 2013. Tiny packet programs for low-latency network control and monitoring. In Proceedings of the Twelfth ACM Workshop on Hot Topics in Networks (HotNets-XII). Association for Computing Machinery, New York, NY, USA, Article 8, 1–7. https://doi.org/10.1145/2535771.2535780
71 Steven Hand, Tim Harris, Alex Ho, Rashid Mehmood; Pervasive Debugging https://www.cl.cam.ac.uk/research/srg/netos/projects/archive/pdb/papers/2004.escience.pdf
72 Cain, B. and D. Towsley. "Generic Multicast Transport Services: Router Support for Multicast Applications." NETWORKING (2000).https://www.semanticscholar.org/paper/Generic-Multicast-Transport-Services%3A-Router-for-Cain-Towsley/aecc9f22d0b0dfd7b416a1837d051bfb92a4a27a

73 Liang Wang, Ben Catterall, Richard Mortier; Probabilistic Synchronous Parallel; 2017 https://arxiv.org/abs/1709.07772

74 Calvert, K.L., Griffioen, J. & Wen, S. Scalable Network Management Using Lightweight Programmable Network Services. J Netw Syst Manage 14, 15–47 (2006). https://doi.org/10.1007/s10922-005-9013-6

75 Yuta Tokusashi, Huynh Tu Dang, Fernando Pedone, Robert Soulé, and Noa Zilberman. 2019. The Case For In-Network Computing On Demand. In Proceedings of the Fourteenth EuroSys Conference 2019 (EuroSys '19). Association for Computing Machinery, New York, NY, USA, Article 21, 1–16. https://doi.org/10.1145/3302424.3303979

76 H. T. Dang et al., "P4xos: Consensus as a Network Service," in IEEE/ACM Transactions on Networking, vol. 28, no. 4, pp. 1726-1738, Aug. 2020, doi: 10.1109/TNET.2020.2992106.

77 Dang, Huynh Tu, Pietro Bressana, Han Wang, Ki-Suh Lee, Noa Zilberman, Hakim Weatherspoon, M. Canini, F. Pedone and R. Soulé. "Partitioned Paxos via the Network Data Plane." ArXiv abs/1901.08806 (2019). https://arxiv.org/abs/1901.08806

78 Zilberman Noa, Moore Andrew W. and Crowcroft Jon A. 2016; From photons to big-data applications: terminating terabits; Phil. Trans. R. Soc. A.3742014044520140445; http://doi.org/10.1098/rsta.2014.0445

79 Swamy, Tushar, Alexander Rucker, M. Shahbaz and K. Olukotun. "Taurus: An Intelligent Data Plane." ArXiv abs/2002.08987 (2020): https://arxiv.org/abs/2002.08987

80 Nicole Hemsoth; Changing Conditions For Neural Network Processing; April 2020 https://www.nextplatform.com/2020/04/07/changing-conditions-for-neural-\network-processing/

81 Whittaker, M., Ailijiang, A., Charapko, A., Demirbas, M., Giridharan, N., Hellerstein, J., Howard, H., Stoica, I., & Szekeres, A. (2020). Scaling Replicated State Machines with Compartmentalization Technical Report. ArXiv, abs/2012.15762 https://arxiv.org/abs/2012.15762

Remote Participants

- Chris Adeniyi-Jones
ARM Ltd. – Cambridge, GB
- Laura Al Wardani
FH Emden, DE
- Uthra Ambalavanan
Robert Bosch GmbH – Renningen, DE
- Gianni Antichi
Queen Mary University of London, GB
- Roberto Bifulco
NEC Laboratories Europe – Heidelberg, DE
- Olivier Bonaventure
UC Louvain, BE
- Kenneth L. Calvert
University of Kentucky – Lexington, US
- Jon Crowcroft
University of Cambridge, GB
- Philip Eardley
BT Applied Research – Ipswich, GB
- T M Rayhan Gias
FH Emden, DE
- Tim Harris
Microsoft Research – Cambridge, GB
- Jianfei He
City University – Hong Kong, HK
- Micchio Honda
University of Edinburgh, GB
- Teemu Kärkkäinen
TU München, DE
- Jussi Kangasharju
University of Helsinki, FI
- Namseok Ko
ETRI – Daejeon, KR
- Michal Król
City – University of London, GB
- Ike Kunze
RWTH Aachen, DE
- Dirk Kutscher
FH Emden, DE
- Julie McCann
Imperial College London, GB
- Jag Minhas
Sensing Feeling – London, GB
- Marie-Jose Montpetit
Concordia University – Montreal, CA
- Naresh Nayak
Robert Bosch GmbH – Stuttgart, DE
- Erik Nordmark
Zededa – San Jose, US
- David Oran
Network Systems Research & Design – Cambridge, US
- Jörg Ott
TU München, DE
- Andy Reid
BT – Ipswich, GB
- Eve M. Schooler
Intel – Santa Clara, US
- Peer Stritzinger
Peer Stritzinger GmbH – Maisach, DE
- Christian Tschudin
Universität Basel, CH
- Klaus Wehrle
RWTH Aachen, DE
- Cedric Westphal
Futurewei – Santa Clara, US
- Peter Willis
BT – Ipswich, GB
- Chenren Xu
Peking University, CN
- Noa Zilberman
University of Oxford, GB

Report from Dagstuhl Seminar 21261

Quantum Complexity: Theory and Application

Edited by

Bill Fefferman[1], Sevag Gharibian[2], Norbert Schuch[3], and Barbara Terhal[4]

1 **University of Chicago, US,** bfefferman@gmail.com
2 **Universität Paderborn, DE,** sevag.gharibian@gmail.com
3 **Universität Wien, AT,** norbert.schuch@gmail.com
4 **TU Delft, NL,** b.m.terhal@tudelft.nl

Abstract

This report documents the program and outcomes of Dagstuhl Seminar 21261 "Quantum Complexity: Theory and Application". The seminar ran from June 27 to July 2 , 2021, and was held in a hybrid format (due to COVID travel restrictions). Of the 55 total participants from 14 countries, 17 participants were on-site, and 38 were remote. Recent advances in both theoretic and experimental aspects of quantum complexity theory were presented and discussed, ranging from new theoretical developments via a "Quantum Strong Exponential Time Hypothesis", to more experimentally oriented talks involving benchmarking of random circuits in quantum supremacy experiments. In addition, an open problem session and a discussion session regarding the current state of the field were included.

Seminar June 27 – July 2, 2021 – http://www.dagstuhl.de/21261

2012 ACM Subject Classification Theory of computation → Quantum complexity theory

Keywords and phrases complexity theory, many-body systems, proof and verification systems, quantum computation, quantum supremacy

Digital Object Identifier 10.4230/DagRep.11.5.76

1 Executive Summary

Bill Fefferman (University of Chicago, US)
Sevag Gharibian (Universität Paderborn, DE)
Norbert Schuch (Universität Wien, AT)
Barbara Terhal (TU Delft, NL)

Background and motivation. Since the seminal discovery of an efficient quantum integer factorization algorithm by Peter Shor in 1994, the field of Quantum Computation has blossomed into a large-scale international effort to build, test, and study the possibilities that information processing using quantum particles may provide. A central role in these developments has been played by Quantum Complexity Theory, a traditionally theoretical realm of research focusing on such questions as: Which physical properties of Nature can be efficiently computed? Can the behavior of an untrusted or noisy quantum computer be verified? What might constitute convincing evidence of "quantum supremacy" over classical computers?

With the first generation of completed "Noisy Intermediate Scale Quantum (NISQ)" experiments already staking quantum supremacy claims, however, the answers to such "traditionally theoretical" questions have taken on an urgent and practical relevance. For example, a complexity theoretic understanding of which realistic physical problems are "just hard enough" for classical computers and "easy enough" for quantum computers is the natural starting point for "quantum supremacy" testbeds. With functioning experimental devices in place, one must next convincingly confirm the device is performing as designed, particularly in the presence of noise. Finally, if the aim of such experiments is to cast doubt on the Extended Church-Turing Thesis, then a strong standard of evidence is required; such a standard must be rigorously stated and developed.

Seminar Topics. This seminar covered a range of topics under the broad umbrella of Quantum Complexity Theory, ranging from highly theoretical to experimentally driven. We briefly overview some of these here; further examples and details are in the included talk abstracts.

Theoretical directions. The field of Quantum Complexity Theory is concerned, broadly speaking, with a rigorous mathematical study of the resources required to perform certain computational tasks. To first order, this involves dividing the "computational world" into two buckets: Easy versus hard problems. However, in reality, the complexity landscape is much finer than this. For example, one might ask – given that problem X has a known efficient quantum algorithm, does there nevertheless exist a *faster* quantum algorithm for X? This typically falls under the classical area of "fine-grained complexity", which has only recently begun to emerge as having a quantum analogue. Conversely, one may ask – is problem X hard only when one wishes to have a high precision answer, becoming easy when a larger margin of error is allowed? Classically, this falls under the umbrella of "hardness of approximation", and which has seen intense study in the guise of the "quantum PCP conjecture". Finally, given that quantum computers are believed more powerful than classical ones, a natural question is: *Do there exist computational problems whose difficulty lies strictly between classical and quantum?* Here, a natural object of study has been so-called "stoquastic" quantum systems, whose time evolution can often be simulated in practice via randomized (i.e. Monte Carlo) techniques, but which nevertheless appear difficult to classically simulate in the *worst case* in a rigorous fashion. Recent advances and the state of the art in all of these topics, as well as a number of others, were discussed at the seminar.

Experimentally motivated directions. The recent explosion of the so-called Noisy Intermediate-Scale Quantum (NISQ) computation era has brought many new questions to the forefront of Quantum Complexity Theory. For example, to date, two of the leading frameworks for experimental demonstration of "quantum supremacy" have been *random circuit sampling* and *Boson sampling*. On the one hand, much progress has been made closing the remaining gaps in the theoretical hardness proofs for these tasks on classical computers. On the other hand, for experiments that *have* been conducted, important practical topics such as how to benchmark such experimental random circuit setups have very recently been studied. Moreover, beyond the quest for quantum supremacy lies the next question: *What practical applications might NISQ devices already prove useful for?* These and related topics were presented and discussed at the seminar.

Participants and program overview. Due to the on-going COVID situation, the seminar was held in hybrid format. This meant that of the 55 total participants joining from 14 countries around the world (from North America to Europe to Asia), 17 were on-site, and

38 were remote. To allow all audience members to participate, a few measures were taken, which arguably worked quite well given the circumstances:

- During each of the seminar's on-site sessions, a Zoom session was projected onto a whiteboard, to which all remote participants were invited. The Zoom participants could see and hear on-site whiteboard and slide presentations, as well as interrupt to ask questions (via the room's loudspeaker system). This made for a reasonably efficient setup in which both on-site and hybrid participants could discuss in real-time. A Slack channel was also set up to ease communication, and by popular request, after talks a virtual Zoom chat room was set up so that the remote participants could also chat amongst themselves.
- To accommodate both types of audience members, a mix of on-site and remote talks were scheduled. On-site talks were typically held in the morning (CEST), allowing remote audience members in Europe Asia to attend. These were held at "standard" times, starting at 9:00 CEST. Remote talks were largely scheduled in the late afternoon and evening (17:00 and 20:00 CEST), this time accessible to North American and European participants.
- Seminar participants Marcel Hinsche (on-site) and James Watson (off-site) graciously offered to act as "technical help volunteers" for local and remote participants, ensuring the hybrid setup ran smoothly for both local and remote attendees.

Regarding the remaining program structure, a strong emphasis was placed on plentiful open time for ad-hoc discussion – typically 14:00 to 17:00 was left open expressly for this purpose. A social outing (hike) was organized by participant Dominik Hangleiter on Wednesday afternoon, and a traditional social night in the music room took place on Wednesday evening.

Acknowledgements. The seminar's participants and organizing committee wholeheartedly thank the Schloss Dagstuhl administrative and technical staff, who before, during, and after the seminar were incredibly supportive, professional, and patient with us quantum computer scientists. Many of the seminars participants, both online and off-line, commented very positively of the experience, citing it as a very welcome break to the stress of the on-going COVID pandemic.

2 Table of Contents

3 Overview of Talks

3.1 An area law for 2D frustration-free spin systems

Anurag Anshu (University of California – Berkeley, US)

Joint work of Anurag Anshu, Itai Arad, David Gosset
Main reference Anurag Anshu, Itai Arad, David Gosset: "An area law for 2D frustration-free spin systems", CoRR, Vol. abs/2103.02492, 2021.
URL https://arxiv.org/abs/2103.02492

We prove that the entanglement entropy of the ground state of a locally gapped frustration-free 2D lattice spin system satisfies an area law with respect to a vertical bipartition of the lattice into left and right regions. We first establish that the ground state projector of any locally gapped frustration-free 1D spin system can be approximated to within error eps by a degree $O(\sqrt{\log 1/eps})$ multivariate polynomial in the interaction terms of the Hamiltonian. This generalizes the optimal bound on the approximate degree of the boolean AND function, which corresponds to the special case of commuting Hamiltonian terms. For 2D spin systems we then construct an approximate ground state projector (AGSP) that employs the optimal 1D approximation in the vicinity of the boundary of the bipartition of interest. This AGSP has sufficiently low entanglement and error to establish the area law using a known technique.

3.2 Quantum fine-grained complexity

Harry Buhrman (CWI – Amsterdam, NL)

Joint work of Harry Buhrman, Bruno Loff, Florian Speelman, and Subhasree Patro
Main reference Harry Buhrman, Bruno Loff, Subhasree Patro, Florian Speelman: "Limits of quantum speed-ups for computational geometry and other problems: Fine-grained complexity via quantum walks", CoRR, Vol. abs/2106.02005, 2021.
URL https://arxiv.org/abs/2106.02005

One of the major challenges in computer science is to establish lower bounds on the resources, usually time, that are needed to solve computational problems. This holds in particular for computational problems that appear in practice. One way towards dealing with this situation is the study of fine-grained complexity where we use special reductions to prove time lower bounds for many diverse problems based on the conjectured hardness of some key problems. For example, computing the edit distance between two strings, a problem that has a practical interest when determining the genetic distance between species based on their DNA, has an algorithm that takes $O(n^2)$ time. Using a fine-grained reduction it can be shown that faster algorithms for edit distance also imply a faster algorithm for the Boolean Satisfiability (SAT) problem (that is believed to not exist). This is evidence that the current edit distance algorithms are optimal. Another problem, besides SAT, that is used as a basis for these reductions is the 3SUM problem. The situation in the quantum regime is no better; almost all known lower bounds for quantum algorithms are defined in terms of query complexity, which doesn't help much for problems for which the best-known algorithms take super-linear time. Therefore, employing fine-grained reductions in the quantum setting seems a natural way forward. However, translating the classical fine-grained reductions directly into the quantum regime is not always possible for various reasons. In this talk, I will present some recent results in which we circumvent these challenges and prove quantum time lower bounds for some problems in BQP conditioned on the conjectured quantum hardness of SAT (and its variants) and the 3SUM problem. This is based on joint work with Bruno Loff, Florian Speelman, and Subhasree Patro.

3.3 Gaussian Boson sampling and its complexity

Abhinav Deshpande (University of Maryland – College Park, US)

Joint work of Abhinav Deshpande, Arthur Mehta, Trevor Vincent, Nicolas Quesada, Marcel Hinsche, Marios Ioannou, Lars Madsen, Jonathan Lavoie, Haoyu Qi, Jens Eisert, Dominik Hangleiter, Bill Fefferman, Ish Dhand
Main reference Abhinav Deshpande, Arthur Mehta, Trevor Vincent, Nicolas Quesada, Marcel Hinsche, Marios Ioannou, Lars Madsen, Jonathan Lavoie, Haoyu Qi, Jens Eisert, Dominik Hangleiter, Bill Fefferman, Ish Dhand: "Quantum Computational Supremacy via High-Dimensional Gaussian Boson Sampling", CoRR, Vol. abs/2102.12474, 2021.
URL https://arxiv.org/abs/2102.12474

Recent demonstrations of a quantum speedup with Gaussian boson sampling have been challenged by new algorithms claiming the absence of this speedup. In this talk, I will discuss the computational hardness of Gaussian boson sampling in an idealized setting of parameters. I will also discuss how, outside of the idealized setting, certain algorithms can simulate some instances of Gaussian boson sampling.

3.4 Linear growth of quantum circuit complexity

Jens Eisert (FU Berlin, DE)

Joint work of Jens Eisert, Jonas Haferkamp, Naga B. T. Kothakonda, Nicole Yunger Halpern, and Philippe Faist
Main reference Jonas Haferkamp, Philippe Faist, Naga B. T. Kothakonda, Jens Eisert, Nicole Yunger Halpern: "Linear growth of quantum circuit complexity", CoRR, Vol. abs/2106.05305, 2021.
URL https://arxiv.org/abs/2106.05305

Quantifying quantum states' complexity is a key problem in various subfields of science, from quantum computing to black-hole physics. We prove a prominent conjecture by Brown and Susskind about how random quantum circuits' complexity increases. Consider constructing a unitary from Haar-random two-qubit quantum gates. Implementing the unitary exactly requires a circuit of some minimal number of gates - the unitary's exact circuit complexity. We prove that this complexity grows linearly in the number of random gates, with unit probability, until saturating after exponentially many random gates. Our proof is surprisingly short, given the established difficulty of lower-bounding the exact circuit complexity. Our strategy combines differential topology and elementary algebraic geometry with an inductive construction of Clifford circuits.

Joint work with Jonas Haferkamp, Philippe Faist, Naga B. T. Kothakonda, and Nicole Yunger Halpern

3.5 Quantum Hardness of Approximation

Lior Eldar (IL)

Main reference Lior Eldar: "Robust Quantum Entanglement at (Nearly) Room Temperature", in Proc. of the 12th Innovations in Theoretical Computer Science Conference, ITCS 2021, January 6-8, 2021, Virtual Conference, LIPIcs, Vol. 185, pp. 49:1–49:20, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2021.
URL http://dx.doi.org/10.4230/LIPIcs.ITCS.2021.49

Quantum entanglement is notoriously hard to maintain and its fragility is arguably the main obstacle preventing us from building a quantum computer. In terms of local Hamiltonians this means that while we know that ground-state of "feasible" quantum systems are highly entangled, we physically can only access the Gibbs states of these quantum systems, and these alas cannot sustain global-scale entanglement.

In this talk we consider the problem of designing systems that exhibit robust quantum entanglement: formally we would like to design a local Hamiltonian for which not only the ground-state is highly entangled but one can also demonstrate that its Gibbs state at non-zero temperature (independent of system size) can only be approximated by deep quantum circuits. Such systems are not known to date.

In [Eldar '21] we show that one can approach such a "holy grail" system and construct a Hamiltonian on n qubits with log-local terms for which the Gibbs state even at nearly-constant temperatures, decaying only at a rate of 1/loglog(n) cannot be approximated by shallow quantum circuits – i.e. of depth less than log(n). The construction involves using state of the art quantum locally testable codes (qLTC), appended with shallow classical decoders for expander codes, together with an analysis of the evolution of thermal errors under qLTCs. The analysis uses the Metropolis Hastings algorithm to show that the errors in the thermal state evolving under a qLTC Hamiltonian tend to form only very sparse errors that are locally correctable – which may be useful elsewhere.

Many open questions remain – among which are improving (reducing) the locality of the construction, and increasing the temperature for which circuit lower bounds can be demonstrated to a constant.

3.6 The power of random quantum circuits

Bill Fefferman (University of Chicago, US)

Joint work of Adam Bouland, Bill Fefferman, Zeph Landau, Yunchao Liu, Umesh Vazirani
Main reference Adam Bouland, Bill Fefferman, Zeph Landau, Yunchao Liu: "Noise and the frontier of quantum supremacy", CoRR, Vol. abs/2102.01738, 2021.
URL https://arxiv.org/abs/2102.01738
Main reference Adam Bouland, Bill Fefferman, Chinmay Nirkhe, Umesh V. Vazirani: ""Quantum Supremacy" and the Complexity of Random Circuit Sampling", in Proc. of the 10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA, LIPIcs, Vol. 124, pp. 15:1–15:2, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, 2019.
URL http://dx.doi.org/10.4230/LIPIcs.ITCS.2019.15

In this talk we will discuss recent results on the power of random quantum circuits, inspired by the "quantum supremacy" experiments of Google and USTC. We will discuss two new results: first we consider the "low noise" scenario in which the goal is to prove the hardness of approximate sampling from the output distribution of a random quantum circuit. The main obstacle faced in prior work on this subject is that the average-case hardness results for computing output probabilities of random circuits are not robust enough to imprecision to connect with the Stockmeyer argument for hardness of sampling. In this work we exponentially improve this robustness to imprecision. In the case of BosonSampling, we bring the proven hardness to within a constant factor in the exponent of the robustness required for hardness of sampling.

Second, we consider the realistic "high noise" scenario. We show that it remains hard to compute the output probabilities of noisy random quantum circuits without error correction, providing the noise rate of the device is below the error detection threshold. This hardness persists despite the fact that these probabilities are exponentially close to uniform. Consequently, the small deviations away from uniformity are hard to compute, formalizing an important intuition behind Google's supremacy claim.

Interestingly, we then argue that these two results are connected, in that any further progress on proving hardness in the "low noise scenario" would require techniques which *do not* work to improve the hardness results in the "high noise scenario".

3.7 (Sub)Exponential advantage of adiabatic quantum computation with no sign problem

András Gilyén (Caltech – Pasadena, US)

Joint work of Hastings, Matthew B.; Gilyén, András; Vazirani, Umesh
Main reference András Gilyén, Matthew B. Hastings, Umesh Vazirani, "(Sub)Exponential advantage of adiabatic quantum computation with no sign problem", Proceedings of the 53rd Annual ACM SIGACT Symposium on Theory of Computing, June 2021, Pages 1357–1369
URL https://doi.org/10.1145/3406325.3451060
Main reference András Gilyén, Matthew B. Hastings, Umesh V. Vazirani: "(Sub)Exponential advantage of adiabatic Quantum computation with no sign problem", in Proc. of the STOC '21: 53rd Annual ACM SIGACT Symposium on Theory of Computing, Virtual Event, Italy, June 21-25, 2021, pp. 1357–1369, ACM, 2021.
URL http://dx.doi.org/10.1145/3406325.3451060

We demonstrate the possibility of (sub)exponential quantum speedup via a quantum algorithm that follows an adiabatic path of a gapped Hamiltonian with no sign problem. This strengthens the superpolynomial separation recently proved by Hastings. The Hamiltonian that exhibits this speed-up comes from the adjacency matrix of an undirected graph, and we can view the adiabatic evolution as an efficient O(poly(n))-time quantum algorithm for finding a specific "EXIT" vertex in the graph given the "ENTRANCE" vertex. On the other hand we show that if the graph is given via an adjacency-list oracle, there is no classical algorithm that finds the "EXIT" with probability greater than $\exp(-n^\delta)$ using at most $\exp(n^\delta)$ queries for $\delta = 1/5 - o(1)$. Our construction of the graph is somewhat similar to the "welded-trees" construction of Childs et al., but uses additional ideas of Hastings for achieving a spectral gap and a short adiabatic path.

3.8 Verifying BQP Computations on Noisy Devices with Minimal Overhead

Elham Kashefi (University of Edinburgh, GB)

Joint work of Elham Kashefi, Dominik Leichtle, Luka Music, Harold Ollivier
Main reference Elham Kashefi, Dominik Leichtle, Luka Music, Harold Ollivier: "Securing Quantum Computations in the NISQ Era", CoRR, Vol. abs/2011.10005, 2021.
Main reference https://arxiv.org/abs/2011.10005

With the development of delegated quantum computation, clients will want to ensure confidentiality of their data and algorithms, and the integrity of their computations. While protocols for blind and verifiable quantum computation exist, they suffer from high overheads and from oversensitivity: When running on noisy devices, imperfections trigger the same detection mechanisms as malicious attacks, resulting in perpetually aborted computations. We introduce the first blind and verifiable protocol for delegating BQP computations to a powerful server with repetition as the only overhead. It is composable and statistically secure with exponentially-low bounds and can tolerate a constant amount of global noise.

3.9 Compact Fermion to Qubit Mappings

Joel David Klassen (Phasecraft – London, GB)

Joint work of Joel Klassen and Charles Derby
Main reference Charles Derby, Joel Klassen, Johannes Bausch, Toby Cubitt: "Compact fermion to qubit mappings", Phys. Rev. B, Vol. 104, p. 035118, American Physical Society, 2021.
URL http://dx.doi.org/10.1103/PhysRevB.104.035118
Main reference Charles Derby, Joel Klassen: "A Compact Fermion to Qubit Mapping Part 2: Alternative Lattice Geometries", CoRR, Vol. abs/2101.10735, 2021.
URL https://arxiv.org/abs/2101.10735

Mappings between fermions and qubits are valuable constructions in physics. To date only a handful exist. In addition to revealing dualities between fermionic and spin systems, such mappings are indispensable in any quantum simulation of fermionic physics on quantum computers. The number of qubits required per fermionic mode, and the locality of mapped fermionic operators strongly impact the cost of such simulations. We present a fermion to qubit mapping that outperforms all previous local mappings in both the qubit to mode ratio and the locality of mapped operators. In addition to these practically useful features, the mapping bears an elegant relationship to the toric code, which we discuss. We additionally discuss the general algebraic framework employed to construct this mapping.

3.10 Provably efficient machine learning for quantum many-body problems

Richard Küng (Johannes Kepler Universität Linz, AT)

Joint work of Richard Küng, Hsin-Yuan Huang, Giacomo Torlai, Victor Albert, John Preskill
Main reference Hsin-Yuan Huang, Richard Kueng, Giacomo Torlai, Victor V. Albert, John Preskill: "Provably efficient machine learning for quantum many-body problems", CoRR, Vol. abs/2106.12627, 2021.
URL https://arxiv.org/abs/2106.12627

Classical machine learning (ML) provides a potentially powerful approach to solving challenging problems in quantum physics and chemistry. However, the advantages of ML over more traditional methods have not been firmly established. We prove that classical ML algorithms can efficiently predict ground state properties of a physical system, after learning from data obtained by measuring related systems. We also prove that classical ML algorithms can efficiently classify a wide range of quantum phases of matter. Our arguments are based on the concept of a classical shadow, a succinct classical description of a quantum state that can be constructed in feasible quantum experiments and be used to predict many properties of the state.

3.11 On QMA Queries with Tree-like Dependencies

Dorian Rudolph (Universität Paderborn, DE)

Joint work of Sevag Gharibian, Dorian Rudolph

The quantum analogue of NP, called QMA (Quantum Merlin Arthur) has the physically motivated complete problem of estimating the ground state energy of a local Hamiltonian. A related problem is simulating the measurement of a local Hamiltonian's ground state. Ambainis (CCC 2014) showed that this problem, denoted APX-SIM (Approximate Simulation), is $\mathrm{P}^{\mathrm{QMA}[\log]}$-complete. $\mathrm{P}^{\mathrm{QMA}[\log]}$ is the class of problems that can be solved by a deterministic polynomial-time Turing machine that may ask a QMA-oracle $\mathrm{O}(\log(n))$ adaptive queries. Gharibian, Piddock, and Yirka (STACS 2020) show that a polynomial number of parallel queries can be simulated using a logarithmic number of adaptive queries and therefore $\mathrm{P}^{\mathrm{QMA}[\log]} = \mathrm{P}^{\|\mathrm{QMA}}$, which also holds for StoqMA.

In the classical setting, an even stronger result is given by Gottlob (JACM 1995): A polynomial number of NP queries with a tree-like dependency graph can be simulated using a logarithmic number of adaptive queries (i.e., $\mathrm{P}^{\mathrm{NP}[\log]} = \mathrm{Trees}(\mathrm{NP})$). More generally, dependent queries can be modeled as a query graph, in which each node contains a query to an oracle for some class C, that is constructed by a uniform circuit taking results from incoming edges as inputs. Within this model, we strengthen Gottlob's result to query graphs with a bounded separator number (this includes bounded treewidth) and apply it to the quantum setting by proving $\mathrm{P}^{\mathrm{C}[\log]} = \mathrm{BSN}(\mathrm{C})$ for C in NP, MA, QCMA, QMA, QMA(2), where BSN(C) denotes the class of problems poly-time reducible to query graph with queries to a C-oracle and a bounded separator number. We further show that query graphs with a logarithmic separator number can be solved by $\mathrm{QP}^{\mathrm{C}[\log^2]}$. We also improve the state of the art for StoqMA by showing that query graphs of constant depth can be solved using a logarithmic number of queries.

3.12 Classical proofs of quantum knowledge

Thomas Vidick (Caltech – Pasadena, US)

Joint work of Thomas Vidick, Tina Zhang
Main reference Thomas Vidick, Tina Zhang: "Classical Proofs of Quantum Knowledge", in Proc. of the Advances in Cryptology – EUROCRYPT 2021, pp. 630–660, Springer International Publishing, 2021.
URL https://doi.org/10.1007/978-3-030-77886-6_22

We define the notion of a proof of knowledge in the setting where the verifier is classical, but the prover is quantum, and where the witness that the prover holds is in general a quantum state. We establish simple properties of our definition, including that, if a nondestructive classical proof of quantum knowledge exists for some state, then that state can be cloned by an unbounded adversary, and that, under certain conditions on the parameters in our definition, a proof of knowledge protocol for a hard-to-clone state can be used as a (destructive) quantum money verification protocol. In addition, we provide two examples of protocols (both inspired by private-key classical verification protocols for quantum money schemes) which we can show to be proofs of quantum knowledge under our definition. Finally, we show that, under our definition, the verification protocol introduced by Mahadev (FOCS 2018) is a classical argument of quantum knowledge for QMA relations.

Participants

- Simon Apers
Free University of Brussels, BE
- Sergio Boixo
Google – Venice, US
- Harry Buhrman
CWI – Amsterdam, NL
- Libor Caha
IBM Research-Zurich, CH
- Jens Eisert
FU Berlin, DE
- Lior Eldar
IL
- Sevag Gharibian
Universität Paderborn, DE
- Dominik Hangleiter
University of Maryland – College Park, US
- Marcel Hinsche
FU Berlin, DE
- Marios Ioannou
FU Berlin, DE
- Robert König
TU München, DE
- Maris Ozols
University of Amsterdam, NL
- Dorian Rudolph
Universität Paderborn, DE
- Norbert Schuch
Universität Wien, AT
- Barbara Terhal
TU Delft, NL
- Frank Verstraete
Ghent University, BE
- Petra Wolf
Universität Trier, DE

Remote Participants

- Scott Aaronson
University of Texas – Austin, US
- Dorit Aharonov
The Hebrew University of Jerusalem, IL
- Anurag Anshu
University of California – Berkeley, US
- Itai Arad
Technion – Haifa, IL
- Johannes Bausch
University of Cambridge, GB
- Adam Bouland
Stanford University, US
- Sergey Bravyi
IBM TJ Watson Research Center – Yorktown Heights, US
- Michael Bremner
University of Technology – Sydney, AU
- Andrew Childs
University of Maryland – College Park, US
- Toby Cubitt
University College London, GB
- Abhinav Deshpande
University of Maryland – College Park, US
- Bill Fefferman
University of Chicago, US
- András Gilyén
Caltech – Pasadena, US
- David Gosset
University of Waterloo, CA
- Alex Grilo
Sorbonne University – Paris, FR
- Sandy Irani
University of California – Irvine, US
- Stacey Jeffery
CWI – Amsterdam, NL
- Elham Kashefi
University of Edinburgh, GB
- Joel David Klassen
Phasecraft – London, GB
- Robert König
TU München, DE
- Robin Kothari
Microsoft Corporation – Redmond, US
- Richard Küng
Johannes Kepler Universität Linz, AT
- Urmila Mahadev
California Institute of Technology – Pasadena, US
- Milad Marvian
University of New Mexico, US
- Reward Mezher
University of Edinburgh, GB
- Tomoyuki Morimae
Kyoto University, JP
- Ramis Movassagh
IBM Research – Boston, US
- Daniel Nagaj
Slovak Academy of Sciences – Bratislava, SK
- Harumichi Nishimura
Nagoya University, JP
- David Pérez García
Complutense University of Madrid, ES
- Stephen Piddock
University of Bristol, GB
- Jamie Sikora
Virginia Polytechnic Institute – Falls Church, US
- Maarten Stroeks
University of Cambridge, GB
- Aarthi Sundaram
MicrosoftMicrosoft Quantum – Redmond, US
- Yuki Takeuchi
NTT – Kyoto, JP
- Ewin Tang
University of Washington – Seattle, US
- Thomas Vidick
Caltech – Pasadena, US
- James Watson
University College London, GB
- Justin Yirka
University of Texas – Austin, US


NOT ALL THOSE WHO WANDER ARE LOST


LZI Karlsruhe
James Watson
Joel Klassen
David Pérez García
Elham Kashefi
Robin Kothari
Justin Yirka
Daniel Nagaj
Alex B. Grilo (he/him)
Ewin Tang
Johannes Bausch
Harry Buhrman
Rawad Mezher
Aarthi Sundaram
Stephen Piddock

Report from Dagstuhl Seminar 21262

Inter-Vehicular Communication – From Edge Support to Vulnerable Road Users

Edited by

Ana Aguiar[1], Onur Altintas[2], Falko Dressler[3], and Gunnar Karlsson[4]

1 **Universidade do Porto, PT,** anaa@fe.up.pt
2 **Toyota Motors North America – Mountain View, US,** onur@us.toyota-itc.com
3 **TU Berlin, DE,** dressler@ccs-labs.org
4 **KTH Royal Institute of Technology – Stockholm, SE,** gk@kth.se

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 21262 "Inter-Vehicular Communication – From Edge Support to Vulnerable Road Users". Looking back at the last decade, one can observe enormous progress in the domain of vehicular networking. In this growing community, many ongoing activities focus on the design of communication protocols to support safety applications, intelligent navigation, and many others. We shifted the focus from basic networking principles to open challenges in edge computing support and, as a novel aspect, on how to integrate so called vulnerable road users (VRU) into the picture.

Seminar June 28, 2021 – http://www.dagstuhl.de/21262
2012 ACM Subject Classification Networks → Network dynamics; Networks → Cyber-physical networks; Networks → Network reliability
Keywords and phrases vehicular networks, vulnerable road users, edge computing, intelligent transportation systems
Digital Object Identifier 10.4230/DagRep.11.5.89

1 Executive Summary

Falko Dressler (TU Berlin, DE)
Ana Aguiar (Universidade do Porto, PT)
Onur Altintas (Toyota Motors North America – Mountain View, US)
Gunnar Karlsson (KTH Royal Institute of Technology – Stockholm, SE)

Looking back at the last decade, one can observe enormous progress in the domain of vehicular networking. In this growing community, many ongoing activities focus on the design of communication protocols to support safety applications, intelligent navigation, cooperative driving and others. Using the terms Vehicular Ad-hoc Networks (VANETs), Inter-Vehicle Communication (IVC), Car-2-X (C2X), or Vehicle-2-X (V2X), many applications – as interesting as challenging – have been envisioned and (at least) partially realized. Very large projects have been initiated to validate the theoretic work in field tests and protocols are being standardized. With the increasing interest from industry, security and privacy have also become crucial aspects in the stage of protocol design in order to support a smooth and carefully planned roll-out. We are now entering an era that might change the game in road traffic management. Many car makers already supply their recent brands with cellular and WiFi modems, some also adding vehicular WLAN (DSRC, ITS-G5) and/or C-V2X technologies, which focus on V2V and V2I communication.

The management and control of network connections among vehicles and between vehicles and an existing network infrastructure is currently one of the most challenging and active research fields in the networking domain. There is a long list of desirable applications that can be grouped into four IVC categories:

1. e-Safety applications that try to make driving safer, e.g., road hazard warning, collision warning;
2. traffic efficiency applications aiming at more efficient and thus greener traffic, e.g., detection of traffic jams, traffic distribution;
3. manufacturer oriented applications, e.g., automatic software updates; and
4. convenience applications, e.g., automatic map updates.

We initiated the "Inter-Vehicular Communication" Dagstuhl Series back in 2010, when a first Dagstuhl Seminar was organized on this topic. The motivation was to bring together experts in this field to investigate the state of the art and to highlight where sufficient solutions already existed. The main outcome of this very inspiring seminar series was that there are indeed areas within this research field where scientific findings are being consolidated and adopted by industry. This was the consensus of quite intriguing discussions among participants from both industry and academia. Yet, even more aspects have been identified where substantial research is still needed.

Some of the findings of the first three seminars in this series have been published not only in the related Dagstuhl reports but also in widely visible magazine articles:

1. Falko Dressler, Frank Kargl, Jörg Ott, Ozan K. Tonguz and Lars Wischhof, "Research Challenges in Inter-Vehicular Communication – Lessons of the 2010 Dagstuhl Seminar," IEEE Communications Magazine, vol. 49 (5), pp. 158-164, May 2011.
2. Falko Dressler, Hannes Hartenstein, Onur Altintas and Ozan K. Tonguz, "Inter-Vehicle Communication – Quo Vadis," IEEE Communications Magazine, vol. 52 (6), pp. 170-177, June 2014.
3. Onur Altintas, Suman Banerjee, Falko Dressler and Geert Heijenk, "Executive Summary – Inter-Vehicular Communication Towards Cooperative Driving," Proceedings of Dagstuhl Seminar 18202 on Inter-Vehicular Communication – Towards Cooperative Driving, Schloss Dagstuhl, Germany, May 2018, pp. 31–59.

Seminars in this series focused on general vehicular communication technologies, security and safety impact, cooperative driving concepts and its implications on communication protocol design, and many more.

We now shifted the focus of this seminar from basic networking principles to open challenges in edge computing support and, as a novel aspect, on how to integrate so called vulnerable road users (VRU) into the picture. Edge computing is currently becoming one of the core building blocks of cellular networks, including 5G, and it is necessary to study how to integrate ICT components of moving systems. The trade-offs of computation distribution, system aspects, and the impact on end-to-end latency are still unanswered. Also, vehicular networking and cooperative driving focuses almost exclusively on cars but leaves out communication and coordination with, for example, pedestrians and bicyclists. For example, many of the existing communication solutions for this scenario were designed without having battery constraints in mind. In the mean-time, some early research has been initiated on this topic, we organized a workshop at INFORMATIK 2019 on VRUs and initial projects report very interesting results on safety features for VRUs. Building upon the great success of the first two seminars, with this follow-up seminar, our goal was to once again bring together experts from all these fields from both academia and industry. The seminar focused intensively on discussions in several working groups. To kick-off these discussions, we invited three keynote talks:

- Distributed machine learning in the vehicle-to-edge continuum by Carla-Fabiana Chiasserini (Polytechnic University of Turin, IT)
- A TechCity Living lab for vehicular-based mobility services in the road by Susana Sargento (Institute of Telecommunications, PT)
- Edge-based increase of awareness and support for all traffic participants by Lars Wolf (TU Braunschweig, DE)

We finally organized the following working groups on some of the most challenging issues related to inter-vehicular communication and cooperative driving:

- Edge Computing: A multi-dimensional techno-economic outlook (i.e., latency, cost, deployment issues, etc.)
- Cooperative Driving Again: Where are we now after 2018? What about vehicular platooning: Is it still alive? What are still unsolved challenges?
- How do we support VRU detection/warning?
- Forget about V2V, we have V2C: Is Vehicle-to-Cloud the Way to Go?

References

1 Falko Dressler, Frank Kargl, Jörg Ott, Ozan K. Tonguz and Lars Wischhof, "Research Challenges in Inter-Vehicular Communication – Lessons of the 2010 Dagstuhl Seminar," IEEE Communications Magazine, vol. 49 (5), pp. 158-164, May 2011.

2 Falko Dressler, Hannes Hartenstein, Onur Altintas and Ozan K. Tonguz, "Inter-Vehicle Communication – Quo Vadis," IEEE Communications Magazine, vol. 52 (6), pp. 170-177, June 2014.

3 Onur Altintas, Suman Banerjee, Falko Dressler and Geert Heijenk, "Executive Summary – Inter-Vehicular Communication Towards Cooperative Driving," Proceedings of Dagstuhl Seminar 18202 on Inter-Vehicular Communication – Towards Cooperative Driving, Schloss Dagstuhl, Germany, May 2018, pp. 31–59.

2 Table of Contents

3 Overview of Talks

3.1 Distributed machine learning in the vehicle-to-edge continuum

Carla-Fabiana Chiasserini (Polytechnic University of Turin, IT)

Main reference Falko Dressler, Carla-Fabiana Chiasserini, Frank H. P. Fitzek, Holger Karl, Renato Lo Cigno, Antonio Capone, Claudio Casetti, Francesco Malandrino, Vincenzo Mancuso, Florian Klingler, Gianluca Rizzo: "V-Edge: Virtual Edge Computing as an Enabler for Novel Microservices and Cooperative Computing", CoRR, Vol. abs/2106.10063, 2021.
URL https://arxiv.org/abs/2106.10063

Intelligent Edge is an emerging paradigm for virtualized mobile services addressing the growing need for AI/ML at the network edge. In this context, mobility safety applications, such as collision avoidance or assisted intersection crossing for vulnerable road users, represent a use case of paramount importance, due to their enormous relevance to society. However, user safety applications requiring ultra-low latency are not the only type of service that needs to be implemented at the edge: networking functions, like virtual radio access networks (vRANs), require to be deployed at the edge as well so as to provide user-to-infrastructure connectivity. The recently proposed Virtual Edge concept looks at such complex ecosystem, and, going beyond Intelligent Edge, aims to find solutions to the challenges that virtualization at the edge poses. In the spirit of Virtual Edge, first we first tackle the dynamic configuration of heterogeneous vRANs providing connectivity to vehicles, as well as to vulnerable road users. Indeed, to fully exploit the potentiality of vRANs in non-stationary environments, an efficient mapping of the rapidly varying context to radio control decisions is not only essential, but also challenging owing to the non-trivial interdependence of network and channel conditions. Our solution, named CAREM, leverages a novel, scalable distributed learning approach, using multiple reinforcement learning agents that allow for dynamic radio resource allocation in the presence of multiple connected users and multiple available links. CAREM meets latency as well as packet loss requirements, enabling the selection of the best radio link among different, available technologies, as well as the modulation and coding scheme (MCS) to be used for packet transmission. It does so by rapidly learning the temporal evolution of the context and associating best actions thereto, based on near real-time feedback on KPI satisfaction as the channel and traffic conditions vary. To demonstrate the benefits of CAREM in real-world settings, we developed a testbed and derived extensive experimental results. Then, giving the scarcity of resources at the edge, we exploit distributed learning to realize the concept of pervasive ML and, so doing, fully leverage the resources offered by the smart city infrastructure and the passing-by vehicles. In particular, drawing on federated learning and distributed stochastic gradient descent, we propose a framework for flexible parallel learning (FPL), achieving both data and model parallelism. Thanks to a newly introduced layer and making an edge node coordinate the learning process, FPL makes distributed learning architectures able to adapt to different network topologies and, hence, flexibly use computing and energy resources in the fog. Further, we investigate how different ways of distributing and parallelizing learning tasks across the participating nodes result in different computation, communication, and energy costs. Our results, obtained using state-of-the-art deep-network architectures and large-scale datasets, show that FPL allows for an excellent trade-off among computational (hence energy) cost, communication overhead, and learning accuracy.

3.2 A TechCity Living lab for vehicular-based mobility services in the road

Susana Sargento (Institute of Telecommunications – Aveiro, PT)

In the framework of the EU project Aveiro STEAM City, researchers in the University of Aveiro and the Institute of Telecommunications, have been deploying an advanced, large-scale communications infrastructure, spread throughout the city of Aveiro, that will be at the service of researchers, digital industries, startups, scale-ups, R&D centers, entrepreneurs and other stakeholders interested in developing, testing or demonstrating concepts, products or services. Supported by state-of-the-art fiber link technology (spread across 16km in the city), reconfigurable radio units, 5G-NR radio and 5G network services, the access infrastructure covers 44 strategic points in the urban area of Aveiro, in the form of smart lamp posts or wall boxes on building facades with communications technologies, edge-based computing units and sensors. The communications infrastructure integrates a communication network with radio terminals, multiprotocol, spread throughout the city, connected by fiber optics to a data processing centre, located at Institute of Telecommunications. Buses and garbage collection vehicles have also been equipped with sensors, which currently record mobility and environmental data, making a complete live map of these parameters in the city, and providing the required data for traffic monitoring and safe driving systems. All these points combine and interconnect a set of sensors, such as mobility sensors (GPS, radars, lidars and video cameras) and environmental sensors (such as temperature, humidity, pollution) with remote data collection units throughout the city, providing enough data to support a wide range of services and applications: from IoT and internet access to citizens, to mobility and intermodal services, smart parking, assisted driving, intelligent transportation systems, environmental monitoring, distribution of information and multimedia content, emergency and safety, health services, among others. This talk dives into the Living Lab and on the vehicular, edge and sensing mechanisms researched to enable a safer mobility environment, with a focus on vulnerable users.

3.3 Edge-based increase of awareness and support for all traffic participants

Lars Wolf (TU Braunschweig, DE)

Research on Inter-Vehicular Communication (IVC) started already some decades ago. Methods to inform others about existence and actual behavior using Cooperative Awareness Message (CAM) and similar have been designed some 20 years ago. During the last decade also means to inform others about observations have been added and standardized with the Collective Perception Message (CPM), thus, increasing the awareness range, e.g., beyond corners. However, these approaches mainly focus on vehicles. Although general ideas and studies to involve Vulnerable Road Users (VRUs) into the overall scenario have been brought up more than a dozen years ago, e.g., in projects such as AMULETT, Car2Ped, and work at NTT Docomo, further research would be useful to integrate VRUs more safely into the traffic landscapes. For that, but also for mere vehicular scenarios, involvement of edge components can help to increase awareness and allow for support of dense traffic situations. Various mobility scenarios can benefit from the inclusion of edge components. However, it brings up new technical needs and concerns. For instance, the edge may and should not only relay messages from vehicles for increased awareness but aggregate several messages to improve scalability, which opens questions about when, what, how to aggregate. Moreover, as it seems now, vehicles may use various communication techniques; the same applies to VRUs as traffic participants with devices worn by humans (smartphones), integrated into protection equipment (helmets, jackets, ...), or attached to bikes/pedelecs, e-scooter, wheelchairs. Overall, real-world scenarios will be complex with many participants, using IVC and non-IVC technologies for communication, many edge components, and needing selection among & handover between them.

Remote Participants

- Ana Aguiar
Universidade do Porto, PT
- Onur Altintas
Toyota Motors North America – Mountain View, US
- Mate Boban
Huawei Technologies – München, DE
- Claudio Casetti
Polytechnic University of Turin, IT
- Wai Chen
China Mobile Research Institute – Beijing, CN
- Carla-Fabiana Chiasserini
Polytechnic University of Turin, IT
- Sinem Coleri
Koc University – Istanbul, TR
- Pedro D'Orey
University of Porto, PT
- Klaus David
Universität Kassel, DE
- Falko Dressler
TU Berlin, DE
- David Eckhoff
TUM CREATE, SG
- Gerhard Fettweis
TU Dresden, DE
- Julian Heinovski
TU Berlin, DE
- Frank Hofmann
Robert Bosch GmbH – Hildesheim, DE
- Gunnar Karlsson
KTH Royal Institute of Technology – Stockholm, SE
- Florian Klingler
Universität Paderborn, DE
- Renato Lo Cigno
University of Brescia, IT
- Ian Robin Marsh
RISE SICS – Stockholm, SE, SE
- Susana Sargento
Institute of Telecommunications – Aveiro, PT
- Michele Segata
Free University of Bozen-Bolzano, IT
- Gürkan Solmaz
NEC Laboratories Europe – Heidelberg, DE
- Christoph Sommer
TU Dresden, DE
- Peter Steenkiste
Carnegie Mellon University – Pittsburgh, US
- Hsin-Mu Tsai
National Taiwan University – Taipei, TW
- Elisabeth Uhlemann
Mälardalen University – Västerås, SE
- Lars Wolf
TU Braunschweig, DE