Report from Dagstuhl Seminar 25312

Building Privacy-Preserving Technologies of Societal Impact

Marina Blanton[*1] **and Liina Kamm**[*2]

1 **University at Buffalo – SUNY, US.** mblanton@buffalo.edu
2 **Cybernetica AS – Tartu, EE.** liina.kamm@cyber.ee

Abstract

This report describes the motivation, purpose, and scope of Dagstuhl Seminar 25312 "Building Privacy-Preserving Technologies of Societal Impact" as well as documents its program and outcomes. This inter-disciplinary seminar brought together computer science researchers and practitioners working on building privacy-enhancing technologies – most notably secure computation applications – and researchers in expertise in other relevant disciplines including law, medicine, and social studies. Besides the applied nature of the seminar that capitalized on the participants' desire to facilitate adoption of privacy-enhancing techniques in real world applications, a unique aspect of this seminar was the shared passion of the participants to use their expertise to build tools for protecting vulnerable populations and for other public good purposes.

Seminar July 27 – August 1, 2025 – https://www.dagstuhl.de/25312
2012 ACM Subject Classification Security and privacy → Cryptography; Security and privacy → Human and societal aspects of security and privacy; Applied computing → Law; Information systems → Data management systems
Keywords and phrases Privacy-enhancing technologies, applications, societal impact, secure computation
Digital Object Identifier 10.4230/DagRep.15.7.280

1 Executive Summary

Marina Blanton (University at Buffalo – SUNY, US, mblanton@buffalo.edu)
Liina Kamm (Cybernetica, EE, liina.kamm@cyber.ee)

License Creative Commons BY 4.0 International license
© Marina Blanton and Liina Kamm

The main goal of the seminar on Building Privacy-Preserving Technologies of Societal Impact was to bring together the main actors in the privacy enhancing technology (PET) sphere to share knowledge of the current state of the art in PETs, get an overview of the most recent cases where these technologies have been used to facilitate data analysis in sensitive social topics, and to discuss ways how PETs could further be exploited to create even more societal impact.

Privacy-preserving techniques such as secure multi-party computation and related areas have matured over the last decades in terms of their speed, accessibility, and usability. However, their propagation into everyday user products continues to be slow. PETs remain largely inaccessible to end users and small or non-profit organizations. A large unexplored potential remains.

Researchers in the community have applied PETs to many application domains and have demonstrated that it is possible to contribute to solving large global challenges such as fighting crime, advancing medical research and patient treatment, strengthening sustainability efforts, reducing gender and race disparities, and much more. In all of this, the fact that certain

[*] Editor / Organizer

Except where otherwise noted, content of this report is licensed under a Creative Commons BY 4.0 International license

information remains confidential is crucial to enabling the functionality which otherwise would not be feasible to carry out. Thus, applied cryptography enables us to achieve what we could not do before using conventional mechanisms and improve individuals' wellbeing.

This seminar's goal was to bring together researchers, whose technical work has contributed to addressing societal challenges, and their allies – and it achieved the goal. The seminar welcomed 28 researchers and practitioners with expertise in computer science, medicine, and law. In their presentations, the participants shared their experiences in building PETs that protect vulnerable populations and contribute to public good as well as their experience with real-world deployment of PETs. The topics included supporting investigative journalists, building privacy-friendly humanitarian aid distribution, privately assessing developmental delays in toddlers, privately matching organ donors, performing medical research, and much more. The presentations also discussed the participants' experience with working with domain experts to design the solutions and deploying these and other applications in the real world. The abstracts of these talks are included in this report.

The participants identified in a brainstorming session a number of open questions they would like to discuss as part of the seminar. The questions ranged from adoption challenges and future strategies for application adoption to non-cryptographer's misconceptions and guidelines for addressing them, to more technical aspects such as interfaces and an evaluation framework, and to legal and funding aspects. The questions were assigned to different discussion sessions, grouping related questions to facilitate coverage of as many topics as possible. The results of these discussions are included in this report. A brief summary is that we need to educate the end users in choosing the right PETs, to convince policy makers of the necessity of PETs, to seek sources of funding for PETs that solve socially sensitive problems, to have the community create and adopt a comparative PET framework for objectively evaluating and comparing different solutions, and to capitalize on successes of privacy applications deployed at a large scale.

This research meeting was very productive and engaging. It had the level of engagement that cannot be found in other settings. And most importantly, it was the passion of the attendees – with technical expertise but deep care for addressing societal challenges – that resulted in the gathering to be hugely successful.

2 Table of Contents

3 Overview of Talks

3.1 Privacy by Design: Supporting Investigative Journalists via PETs

Kasra EdalatNejad (TU Darmstadt, DE)

License Creative Commons BY 4.0 International license
© Kasra EdalatNejad

Investigative journalists collect large numbers of digital documents during their investigations. These documents can greatly benefit other journalists' work. However, many of these documents contain sensitive information. Hence, possessing such documents can endanger reporters, their stories, and their sources. Consequently, many documents are used only for single, local, investigations. We presented DatashareNetwork, a decentralized and privacy-preserving search system that enables journalists worldwide to find documents via a dedicated network of peers, as the first search engine designed by journalists for journalists in 2020 to address this problem. We start the talk by introducing real-world problems that investigative journalists face and describe DatashareNetwork as a possible solution. Then, we discuss the practical challenges of moving forward from an academic prototype to deploying DatashareNetwork for the International Consortium of Investigative (ICIJ). This talk covers (1) our joint requirement gathering and (2) design with journalists, (3) a user study to help ICIJ with presenting the privacy property of our system to journalists and making utility/privacy trade-off decisions, (4) deployment challenges to integrate DatashareNetwork into ICIJ's IT infrastructure, and finally (5) open problems that require more attention from the community.

3.2 Private Set Intersection for the Society

Thomas Schneider (TU Darmstadt, DE)

License Creative Commons BY 4.0 International license
© Thomas Schneider
Joint work of Thomas Schneider, members of the ENCRYPTO group & collaborators

Private set intersection allows two parties to privately compute the intersection of their inputs or even variants of this functionality. The first PSI protocol is due to C. A. Meadows (IEEE S&P'86) based on Diffie-Hellman and since then a lot of fruitful research has happened in this area, making it more and more practical. In this talk, I will summarize some of our works that investigate applications of PSI of societal impact.

In the first part, based on Kales et al. (USENIX Security'19) and Hagen et al. (NDSS'21), we look at contact discovery, which is a feature in almost all mobile messenger apps. Here, the goal is to match a few 1000 address book contacts against potentially billions of users of that service to determine which of the contacts are using the service. We show that the approach of hashing phone numbers as used by many messengers provides literally no additional privacy because phone numbers are structured and have little entropy, so they can be inverted within milliseconds. Protocols for unbalanced PSI provide the required privacy guarantees, but still, communication is a major bottleneck unless resorting to the multi-server setting, as we showed in Hetz et al. (ESORICS'23). Moreover, we show that the contact discovery functionality as implemented in the messengers can be exploited to mass-collect highly sensitive information like profile pictures or online status for a large number of phone numbers, and after our responsible disclosure the providers adjusted their rate limits.

The second part is based on Treiber et al. (WPES'22). This is an interdisciplinary effort together with law experts from Goethe University Frankfurt and the Police Academy Hamburg, where we investigate how law enforcement agencies can selectively and privately exchange data from both a legal and a technical perspective. We propose a system for lawful information exchange between law enforcement agencies using MPC and private set intersection, which has much higher privacy guarantees than the currently discussed approaches like pooling data in the clear in a data warehouse or exchanging hashed identities.

3.3 Privacy-Preserving Humanitarian Aid Distribution

Wouter Lueks (CISPA – Saarbrücken, DE)

License Creative Commons BY 4.0 International license
© Wouter Lueks

Humanitarian aid-distribution programs help bring physical goods to people in need. Traditional paper-based solutions to support aid distribution do not scale to large populations and are hard to secure. Existing digital solutions, on the other hand, risk creating serious harms to recipients by collecting large amount of personal information including sometimes biometric data.

This talk covers our collaboration with the International Committee of the Red Cross to design digital aid-distribution systems with strong protection. We'll see how we used privacy-enhancing technologies to prevent harm to vulnerable aid recipients; and touch on our experience working with humanitarian organizations.

3.4 Regional Risk Monitoring of Developmental Delay in Toddlers Using MPC: Use Case Overview, and Aspects of Getting Secret-Shared Inputs via the Browser

Niek Bouman (Roseman Labs – Utrecht, NL)

License Creative Commons BY 4.0 International license
© Niek Bouman

The first part of this talk describes a real-world use case of one of Roseman Labs' clients, the Municipality of Rotterdam (the Netherlands), about monitoring the effectiveness of municipality-level policy concerning developmental delay in toddlers. You will learn how MPC was employed to achieve GDPR-compliance and how its use shortened the feedback cycle from 1,5 year to one month; creating real impact because the toddlers directly benefit from this.

The second, more technical part is about getting input data for MPC (in the "outsourcing model") from users via a web application, and discusses the problem of ensuring integrity of the web app, e.g., in a scenario where the web server gets compromised.

3.5 Addressing the Kidney Exchange Problem – From Theory Towards Practice

Susanne Wetzel (Stevens Institute of Technology – Hoboken, US)

License Creative Commons BY 4.0 International license
© Susanne Wetzel

Today, thousands of patients in Germany alone suffer from severe kidney disease and are on the waiting list to receive a kidney transplant. Receiving a living kidney donation is an alternative to waiting for a post-mortom kidney transplant. The main challenge with living kidney donation, however, is to find organ donors who are medically compatible with the patients. Also, not all countries allow kidney exchange due to the fear of manipulation, corruption, and coercion. Yet, in many countries where kidney exchange is allowed, there already are centralized platforms to facilitate the exchange – which in turn raises security and privacy concerns.

Our work seeks to develop privacy preserving protocols for solving the kidney exchange problem aiming to address these concerns.

In this talk, we discussed the challenges we encountered in carrying out this international and interdisciplinary project.

This is joint work with the research group of Professor Ulrike Meyer at RWTH Aachen University.

3.6 Secure Multiparty Computation in a European Clinical Study

Hendrik Ballhausen (LMU – München, DE)

License Creative Commons BY 4.0 International license
© Hendrik Ballhausen
Main reference Hendrik Ballhausen, Stefanie Corradini, Claus Belka, Dan Bogdanov, Luca Boldrini, Francesco Bono, Christian Goelz, Guillaume Landry, Giulia Panza, Katia Parodi, Riivo Talviste, Huong Elena Tran, Maria Antonietta Gambacorta, Sebastian Marschner: "Privacy-friendly evaluation of patient data with secure multiparty computation in a European pilot study", npj Digit. Medicine, Vol. 7(1), 2024.
URL http://dx.doi.org/10.1038/S41746-024-01293-4

In multicentric studies, data sharing between institutions might negatively impact patient privacy or data security. An alternative is federated analysis by secure multiparty computation. This pilot study demonstrates an architecture and implementation addressing both technical challenges and legal difficulties in the particularly demanding setting of clinical research on cancer patients within the strict European regulation on patient privacy and data protection.

3.7 Screening DNA Synthesis Orders for Hazards: Efficiently and with Privacy

Carsten Baum (Technical University of Denmark – Lyngby, DK)

License Creative Commons BY 4.0 International license
© Carsten Baum

DNA synthesis has become a ubiquitous tool in biological and medical research. It allows to synthesize arbitrarily built DNA strands based on digital descriptions only. However, DNA synthesis can be used for nefarious purposes.

The SecureDNA project built an efficient software tool to screen DNA orders for hazardous sequences with some privacy guarantees. In this talk, I describe the general architecture, the development process and some cryptographic underpinnings.

3.8 The German Self-Determination Act: MPC as a Tool for Better Privacy in Administrative Data Exchange?

Andreas Brüggemann (TU Darmstadt, DE)

License Creative Commons BY 4.0 International license
© Andreas Brüggemann
Joint work of Linda Seyda, Andreas Brüggemann, Gerrit Hornung, Thomas Schneider
Main reference Linda Seyda, Andreas Brüggemann, Gerrit Hornung, Thomas Schneider: "Multi-Party Computation als Instrument zur Umsetzung datenschutzkonformer behördlicher Datenabgleiche: Eine interdisziplinäre Analyse am Beispiel der Diskussionen um das Gesetz zur Selbstbestimmung über den Geschlechtseintrag", in Proc. of the 54. Jahrestagung der Gesellschaft für Informatik, INFORMATIK 2024 - Lock in or log out? Wie digitale Souveränität gelingt, Wiesbaden, Germany, September 24-26, 2024, LNI, Vol. P-352, pp. 153–167, Gesellschaft für Informatik, Bonn, 2024.
URL http://dx.doi.org/10.18420/INF2024_11

In 2024, the German self-determination act (SBGG) was passed, enabling transgender, intersex, and non-binary people to change their name and gender entries by self-declaration. The original draft of the SBGG included that all entry changes are automatically forwarded to many law enforcement agencies to allow them to keep their internal databases up to date. This data transfer would have been unprecedented with name changes due to other reasons not being transferred to law enforcement without cause. Furthermore, it would have amplified the risk of lists of transgender, intersex, and non-binary people being created in a situation where German police has already been reported to collect lists of queer people in the past while also, violence against queer people has recently been increasing.

In an interdisciplinary collaboration on the intersection of law and cryptography, we have analyzed the risks of the planned data transfer as well as the underlying considerations from the perspective of data protection. Not only is the proportionality of the transfer generally at least highly questionable, but the transfer also contains data about individuals completely unknown to law enforcement where a legitimate interest on the side of law enforcement cannot apply. Yet, cryptographic means such as private set intersection (PSI) have been ignored in the lawmaking process while they would have enabled to at least protect information about most affected people. We provide an analysis of how PSI could have improved parts of the problem, which problems it could not have resolved, and which smaller, but new risks it might have introduced. While the data transfer has eventually been removed before the law was passed, this was only done for consistency reasons with the plan to reintroduce a similar mechanism for all name changes in Germany while an even worse version is now being discussed specifically for the SBGG. Our research in this lawmaking process showcases how concrete risks and data protection concerns were ignored while generally, privacy enhancing technologies such as PSI appear not to be considered or even known to the responsible lawmakers while they would enable possible compromises.

3.9 Verifiable Carbon Accounting in Supply Chains

Jonathan Heiß (TU Berlin, DE)

License Creative Commons BY 4.0 International license
© Jonathan Heiß
Main reference Jonathan Heiss, Tahir Oegel, Mehran Shakeri, Stefan Tai: "Verifiable Carbon Accounting in Supply Chains", IEEE Trans. Serv. Comput., Vol. 17(4), pp. 1861–1874, 2024.
URL http://dx.doi.org/10.1109/TSC.2023.3332831

Trustworthy data sharing in carbon accounting is hindered by confidentiality constraints that prevent consumers from understanding how providers construct emission data. This opacity enables greenwashing and undermines consistency across value chains. Zero-knowledge proofs (SNARKs) combined with verifiable data structures, and other privacy-enhancing technologies (PETs) address these limitations by enabling verification without revealing sensitive business internals. This talk introduces verifiable carbon accounting (VCA) as a framework synthesizing multiple PETs for scalable, confidentiality-preserving emission data verification and discusses adoption strategies to integrate VCA into existing accounting standards as an alternative to costly, non-scalable third-party verification.

3.10 MPC in Practice: Lessons from Germany's Public Sector and Future Applications in Sustainability

Ágnes Kiss (SINE Foundation – Berlin, DE)

License Creative Commons BY 4.0 International license
© Ágnes Kiss

This presentation explores the practical implementation of secure Multi-Party Computation (MPC) through real-world use cases in Germany's public sector and emerging applications in sustainability.

Our work with German municipalities reveals that successful MPC deployment requires more than technical solutions. While trust between parties was not the primary barrier, data privacy concerns and GDPR compliance proved crucial. We present concrete use cases including disaster management, childhood health examination analytics, and vaccination verification, highlighting the importance of finding committed stakeholders willing to pioneer new approaches. The technical foundation relies on our MPC framework combining Garble (high-level Rust-like programming for Boolean circuits) and polytune (full-threshold MPC with authenticated garbling).

Looking toward sustainability applications, we identify significant potential for privacy-preserving technologies in carbon footprint verification and supply chain transparency. Product Carbon Footprint (PCF) calculations require verifiable accuracy and transparency, yet involve sensitive trade secrets across multiple stakeholders. MPC offers a path to shared insights from primary data without compromising competitive advantages, enabling e.g., statistical outlier detection, PCF categorization, and input validation while preserving data confidentiality. The potential of Privacy-Enhancing Technologies in sustainability contexts remains largely unexplored, presenting opportunities for interdisciplinary research and practical applications.

3.11 Secure Analytics with MPC Made Practical: Lessons Learned from Medical and Workforce Data Use Cases

John Liagouris (Boston University, US)

License © Creative Commons BY 4.0 International license
© John Liagouris

"The performance of MPC-based approaches is so low that practical applicability is not in sight." This is a review excerpt of a paper I co-authored, describing our vision to use multiparty computation (MPC) for secure data analytics in the cloud. In this talk, I will share how – 4+ years later – we have realized this unlikely vision and more. I will first explain the legitimate skepticism of the particular reviewer and why past results indicated that MPC protocols were impractical for complex analytics. I will then argue that careful system design and cross-layer optimizations can not only amortize MPC costs, but also achieve scalability to large inputs and complex workloads, without compromising security. I will present the BU Secure Analytics Stack, our unified software architecture for secure collaborative data analysis, and discuss some lessons learned from real use cases. Finally, I will show performance results for secure relational and time series analytics at a scale that a few years ago was only possible with information leakage or the use of trusted compute.

3.12 MPC for Securing Auctions in a P2P Electricity Market

Mariana Gama (KU Leuven, BE)

License © Creative Commons BY 4.0 International license
© Mariana Gama

The development of the smart grid and introduction of smart metering devices facilitates the emergence of new applications within the energy domain, among which is the p2p electricity trading market. This market would allow owners of renewable energy sources to sell any excess electricity generated directly to other users, incentivising users to acquire renewable energy sources and diminishing transmission losses by promoting electricity exchanges among neighbours. However, there are several privacy concerns associated with the p2p electricity market, as exposing fine-grained metering data exposes sensitive information about users' habits.

This talk introduces MPC-based auction mechanisms for the future p2p electricity market. We discuss the possibility of prioritising both low-volume orders and close neighbours when matching orders for intraday p2p electricity trading, and propose a day-ahead flexibility market where users' consumption schedule is optimised with respect to day-ahead electricity prices.

3.13 PETs in Practice

Liina Kamm (Cybernetica AS – Tartu, EE)

License Creative Commons BY 4.0 International license
© Liina Kamm

Privacy enhancing technologies (PETs) have been around for decades. Their feasibility has been an area of research for tens of research teams. However the practical uptake of the more complex and privacy-preserving solutions has been slow.

In the beginning of 2023, Estonia conducted a research project on privacy enhancing technologies to work out a concept and roadmap for deploying these technologies in e-government. We interviewed people from 18 state agencies to find out their expectations and requirements for the use of different PETs. This lead us to compile a PET concept that gives and overview of the technologies and describes the generalised usage archetypes for e-government. The PET reoadmap gives a concrete way forward.

We discussed the lessons learned from this work and distilled from this a list of incentives and barriers that possible users and customers see.

3.14 Securing Satellite Rendezvous and Proximity Operations with MPC

Kevin Butler (University of Florida – Gainesville, US)

License Creative Commons BY 4.0 International license
© Kevin Butler

Space is emerging as a contested and congested environment, driven by the accelerating deployment of satellites from government and commercial organizations. Consequently, in-space security and privacy have become a significant concern, particularly related to satellite rendezvous and proximity operations (RPO). This talk describes our efforts to develop and implement critical RPO algorithms on radiation-tolerant hardware suitable for deployment in low Earth orbit using MPC to identify and ensure privacy of appropriate inputs, implemented with the MP-SPDZ framework. We highlight the need for optimizations and careful mission considerations when deploying MPC in this specialized environment. From a social impact perspective, we also discuss the potential for MPC to address interpersonal threats posed by continuous location-sharing services.

3.15 Scaling Privacy: Cloud-Native MPC, Data Anonymization, and the Path to Open Linux Ecosystem

Hossein Yalame (Robert Bosch GmbH – Renningen, DE)

License Creative Commons BY 4.0 International license
© Hossein Yalame

As organizations increasingly rely on data-driven decision-making, preserving privacy at scale has become a critical challenge. This work presents CarbyneStack, a cloud-native framework for secure multi-party computation (MPC), developed and deployed at Bosch to enable privacy-preserving analytics across global subsidiaries. Bosch has contributed CarbyneStack

to the Linux Foundation Europe, fostering an open ecosystem for privacy-enhancing technologies. We discuss how this open-source initiative promotes interoperability, transparency, and community-driven innovation while maintaining enterprise-grade performance. Our results highlight that having MPC in a cloud-native architecture not only scales privacy-preserving analytics but also accelerates the adoption of secure, privacy-focused solutions across industries.

3.16 Increasing the Impact of an MPC (or Any PET) Application

Dan Bogdanov (Cybernetica AS – Tartu, EE)

License Creative Commons BY 4.0 International license
© Dan Bogdanov

What is needed to grow the impact of a Privacy Enhancing Technology system?

Easier to achieve:

- People with skills – PET and MPC skill dissemination works well
- Open code – we found 56 projects/products (not all open, but there are more)
- Funding for pilots – EU and US have both invested > 300M€/$ into PETs by now!
- Regulatory support – it is doable for non-PET solutions, so it's a matter of will

Harder to achieve:

- Sponsor/customer deciding to use PETs – fears of cost, alternative solutions
- Working code – keeping cryptography code running for years is non-trivial
- Ongoing operations – need to find money and skills to run the (distributed!) system
- Trust building – need to convince end users and the public that this will be secure

The talk focuses on progress in the hard-to-achieve aspects of PET growth.

3.17 Six Years of MP-SPDZ in the Wild

Marcel Keller (CSIRO – Eveleigh, AU)

License Creative Commons BY 4.0 International license
© Marcel Keller

In the few years since publication, MP-SPDZ has received more than 1000 GitHub issues and hundreds of citations. I will present a categorization of issues and citations that mention the usage of the framework. I will also enumerate the most frequent themes in the issues and give my view on how to address them, whether it's worth it, or possible at all. This is to gather views from the community and foster discussion.

3.18 Getting up to (MP-)SPDZ – Challenges Faced by Developers Getting into MPC

Vincent Ehrmanntraut (RWTH Aachen, DE)

License Creative Commons BY 4.0 International license
© Vincent Ehrmanntraut

For most people, MPC is a dauntig subject to approach at first, as a master of algorithmic concepts and implementation details is needed to design effective protocols. This results in a very steep learning curve.

In this talk, I share insights from advising bachelor and master theses to point out specific pain points and present some ideas that might help to flatten the learing curve. The talk then pivots to make a case that the whole field of (high-level) MPC protocols needs to establish evaluation practices that improve the comparability of protocols across papers.

3.19 Challenges of Making PETs with Societal Impact Usable & Trustworthy

Simone Fischer-Hübner (Karlstad University, SE)

License Creative Commons BY 4.0 International license
© Simone Fischer-Hübner

Privacy Enhancing Technologies (PETs) will only be successfully deployed if they are usable and trustworthy. However, for making PETs usable, various challenges need to be addressed: First of all, there are no "one size fits all" solutions. The context of a data processing applications and demographics of users, and in particular their cultural and gender backgrounds, need to be considered for the usable design of PETs. Secondly, explaining PETs that are based on "crypto magic" operations constitutes a challenge and additionally technical background knowledge may have a negative impact on the users' mental models. Thirdly, for the configuration of PETs, interdisciplinary expertise may be needed and complex trade-offs between protection goals may have to be made. In this talk, these usability challenges are illustrated with six user studies on PETs that can have societal impact, including a Selective Authentic EHR Exchange Service based on malleable signatures, privacy preserving data analytics based on homomorphic encryption, differential privacy and functional encryption, and lastly secret sharing for secure cloud storage. The presentation concludes with guidelines and recommendations for usable and trustworthy PETs.

3.20 How People would Trust a Government based on its Choice of a Digital Identity Wallet?

Kazue Sako (Waseda University – Tokyo, JP)

License Creative Commons BY 4.0 International license
© Kazue Sako

I would like to share preliminary results from a survey we conducted in five countries, each with 800 participants, exploring how people would trust a hypothetical government based on its choice of a digital identity wallet – especially when the wallet supports unlinkable selective disclosure. The study was conducted with a social psychology researcher, who carefully designed the survey and handled the statistical analysis.

3.21 MPC and PETs from a GDPR Perspective

Meiko Jensen (Karlstad University, SE)

License Creative Commons BY 4.0 International license
© Meiko Jensen

Multiparty Computations are commonly perceived as a privacy-enhancing technology of high quality. However, when it comes to compliance to data protection laws like the European GDPR, the implementation of MPC schemes may have unintended side effects. In this talk we discussed different aspects of the GDPR and their relation to PETs in general and MPC specifically. Here, we considered three cases: MPC as an anonymization tool, MPC as a risk reduction tool according to Art. 35 GDPR, and MPC as a methodology to implement Data Protection by Design (Art. 25) and Security of Processing (Art. 32).

For anonymization, we noted that the intermediate results of an MPC computation may still contain linkable information of a data subject, hence, for most cases, applying MPC does not result in a sufficient level of anonymization as defined in the GDPR. Hence, MPC should not be considered as an anonymization tool in order to get out of the scope of the GDPR.

With respect to risk reduction, we noticed that the implementation of MPC does in fact help with confidentiality and integrity risk mitigation, but its implementation opens up several other risk vectors, such as availability issues (due to performance and network workloads), administration issues, and issues related to the addition of new stakeholders (the compute nodes that may belong to different organizations, hence causing either joint controllership or controller-processor relationships in the sense of GDPR). Hence, applying MPC is a trade-off of risk mitigation and novel risk addition, which should be evaluated carefully.

Finally, though MPC may clearly be seen as a sound approach to implement privacy by design and security of processing, it nevertheless is not mandatory to be implemented. Other PETs of lower quality may suffice to cover these GDPR requirements, so the incentive to utilize MPC, given its novel risk additions and pitfalls, must be evaluated carefully.

Based on these observations, we discussed different approaches to improve the status of MPC under GDPR considerations. We identified that it needs more success stories of MPC implementation, to get out of the "PET graveyard" of piloted, but never marketed MPC tools. Also, an update to GDPR may help catering for the specific needs of MPC adoption. Finally, we identified that MPC is one of the very few privacy-enhancing technologies that is capable of protecting against the "trump attacker model," where a government turns rogue against its own institutions and tries to misuse existing data, including the removal or bypass of protection mechanisms (e.g., by forcing decryption and secret key revelation). If one or more parties of the MPC implementation are out of scope of such an attacker, e.g., in a different jurisdiction, MPC could withstand such an attack incident.

3.22 Coordinating MPC Parties with a Byzantine Fault-Safe Distributed Database

Mark Abspoel (Roseman Labs – Utrecht, NL)

License Creative Commons BY 4.0 International license
© Mark Abspoel

Consider an outsourced secure multiparty computation scenario with a small number of long-running computation parties that can execute multiple computations, as follows. Input parties asynchronously secret-share confidential data and distribute the shares to the computation parties. An output party can request for a computation to be run on the secret-shared input data, subject to manual approval by a fixed set of approver users.

This talk is about how a distributed database of computation-related metadata can be facilitated by the computation parties, that may contain the schemas of secret-shared data, the approvals that are given by the approver users, audit logs of computations, et cetera. The database should be resilient against an active adversary corrupting a dishonest majority of computation parties.

We develop a novel Byzantine-fault-safe distributed database through a synchronizing layer on top of existing relational databases, that handles access control and ensures determinism. It is based on a novel protocol for ledger consensus with abort, where we circumvent the impossibility result for ledger consensus with a dishonest majority by replacing liveness by a weaker notion of liveness with abort.

3.23 Privacy-Preserving Collaborative Learning

Sinem Sav (Bilkent University – Ankara, TR)

License Creative Commons BY 4.0 International license
© Sinem Sav

In this talk, I'll share my experience through the evolving space of privacy-preserving collaborative learning. I'll reflect on the challenges we've faced in implementing the first federated and privacy-preserving machine learning framework and its application to the biomedical domain such as: (1) Conflicting application needs, (2) Inconsistent notions of privacy across domains, and (3) The growing tension between theory and practice.

Key discussion points will include: (1) Whether tailored, application-specific solutions are more effective than general-purpose ones; (2) How we might reconcile varying interpretations of "privacy-preservation" across stakeholders; and (3) What it will take to make techniques like homomorphic encryption or differential privacy more practical. I'll also mention the details that often get overlooked, like how to set up a collaborative learning pipeline, e.g., from hyperparameter tuning or data normalization.

3.24 They Drop, We Chat: Steganographic Censorship Circumvention for Chat-based Applications

Boya Wang (EPFL – Lausanne, CH)

License Creative Commons BY 4.0 International license
© Boya Wang

Censorship prevents communications that are against the interest of the censor, and hence, silences the conversations key to a healthy society. Chat applications become a major target of censorship as they become the most popular tools for our communication nowadays. One type of censorship is sensitive-word filtering (SWF) which exists in popular non-E2EE chat applications such as WeiXin/WeChat. The censor deploys an adaptive blocklist of words at the application server to drop matched messages in real-time. In this work, we investigate steganographic circumvention system of content-based censorship in chat applications. Previous work studies the security of steganographic circumvention systems in isolation without considering system-wise or contextual requirements, which results for a mismatch between theoretical security guarantee and practical undetectability. We fill this gap by proposing a new design of the circumvention system, provide two steganographic instantiations, and conduct an empirical evaluation.

3.25 Understanding Differential Privacy Adoption Challenges for Technical Implementers

Rebecca Wright (Barnard College, Columbia University – New York, US)

License Creative Commons BY 4.0 International license
© Rebecca Wright
Joint work of Liudas Panavas, Saeyoung Rho, Hari Bhimaraju, Wynee Pintado, Rebecca Wright, Rachel Cummings

Introduced by Dwork et al. in 2006, differential privacy can provide rigorous mathematical guarantees of user privacy. Differential privacy has been a research success and has been adopted by a number of large organizations, including companies like Apple, Google, and Microsoft, as well as the US Census Bureau and the Israeli Ministry of Health. Nonetheless, there are socio-technical challenges that have limited broad adoption.

In this talk, we describe our interview-based research study seeking specifically to address the question: For technical individuals newly introduced to differential privacy, which concepts are most challenging to grasp, and what factors contribute to these challenges? We interviewed 10 subjects, people knowledgeable about DP who have managed or worked with software engineers not previously knowledgeable about DP to implement it.

We discuss the four primary themes we find in the interviews: (1) scoping as a structural barrier, (2) the need for expert judgment, (3) tunable vs. codified parameters: as simple as possible, but not oversimplified, and (4) visualizations facilitate communication. We conclude with some behind-the-scenes stories of our project as well as possible next steps.

4 Discussion Sessions

4.1 Where Does MPC Have an Actual Advantage?

The topic of this discussion session was to determine where secure multi-party computation has advantage compared to other technologies and thus where it can see adoption success.

We discussed the following question: Assume trusted execution environments (TEEs) are secure, and that fully homomorphic encryption (FHE) is fast. Is there still a point to secure multiparty computation (MPC)?

We came up with three main advantages of MPC.

First, since MPC is software-based, anyone can inspect the source code, or even build an implementation based on a specification. This is in contrast to hardware-based solutions such as TEEs or FHE accelerators, where only some countries or organizations may have the capabilities to produce and/or inspect them. This point can be important due to geopolitical considerations, e.g., where there are different countries that need a symmetry of power.

Second, due to its distributed nature, MPC offers more control over which computations are executed on confidential data. For example, parties that supply data can choose to take on a role as a computation party, which enforces the need for its consent and cooperation (given a suitable security model).

Third, it is easier to fix vulnerabilities in MPC due to its software-based nature, which is especially important for long-running systems. By contrast, fixing vulnerabilities in hardware-based solutions can be costly and hard to achieve.

We also discussed a few aspects that hinder adoption of MPC. For lawyers, TEEs and MPC are relatively indistinguishable. This also applies to policy makers: because MPC is hard to understand, it does not really make its way into policies and regulations. For cloud providers, promoting MPC is risky, because it may lead users to infer that the cloud is not as secure as they claim. Consultancy companies may be aware of MPC, but it can also be difficult to come up with concrete use cases that generate business value.

As a side remark, it is interesting that MPC still has a place in key management, or is used as an alternative for a hardware security module, for applications with high security needs (e.g., cryptocurrency wallets). Perhaps some users do think MPC is more secure than hardware-based solutions.

4.2 How Should We Speak to Non-Cryptographers and How Can We Convince Medical Data Guardians?

Below is the summary of the discussion on guidelines on talking to non-crypto persons and developing a roadmap to talk to medical data guardians (who, in our experience, resist the use of MPC).

The most prominent theme of the discussion was that talking to non-crypto persons requires knowledge of the non-crypto domain, and a good understanding of the non-crypto person's pain points / problems.

Furthermore, demos in the context of the non-crypto person and general success stories, where MPC already is used successfully, are helpful to convince the non-crypto person that using MPC might be a good idea. On the other hand, telling the non-crypto person that there is a privacy problem, and then attempting to sell the non-crypto person on that privacy problem usually fails. The main exception is when the non-crypto person is already passionate about that privacy problem.

On a conversational level, non-crypto persons generally do not like the word attacker, or being called untrustworthy. A possible remedy is to introduce the "Trump Attacker," i.e., arguing that the successor of the non-crypto person might not be trustworthy.

Simplification of cryptographic concepts is another important aspect. Non-crypto persons usually need harder simplifications than initially assumed by cryptographers. Also, cryptographers tend to overly focus on the negatives of their technologies, e.g., by starting with the assumptions their systems require or directly pointing out possible attacks or implementation flaws. While it is important to remain realistic and not to oversell MPC, it may be necessary to initially mask uncertainty and have (faux) confidence in proposed solutions.

The discussion on healthcare-specific topics was very brief, and identified two problems: First, the medical persons might not see the need to go beyond legal compliance, e.g., with HIPPA. Secondly, hospitals (especially in the EU) already struggle with IT, and thus need extra convincing that adopting MPC won't burden their IT too much. Therefore, it might be necessary to "bring your own IT staff" for pilot projects.

4.3 What Misconceptions about Privacy-Enhancing Technologies Do Users Have?

We talked about what misconceptions users typically have about privacy technologies, and how to address some of them.

Many users are not able to distinguish well between different security technologies. For example, users may believe that privacy-enhancing technologies (PETs) like secure multiparty computation work akin more familiar cryptographic tools such as encryption. Or they think that end-to-end encryption also protects metadata. We have also seen that some users think that "private browsing" is a tool for connection security.

To address these misunderstandings, it is helpful to explain first what privacy guarantees are provided by a system, and to leave details about the privacy technology as secondary information.

For the privacy guarantees, we could come up with a more standardized way to, for a given system, communicate what data is accessible to what party, and under what guarantees (e.g., for WhatsApp, message content is accessible to a user's phone and not to the server, due to end-to-end encryption; but metadata is accessible to both the user's phone and the server). One question is who will write these: are creators of the systems going to do this, at the risk of presenting their system in a less favorable way, or is this best left to independent parties such as consumer organizations?

To explain PETs, we can make the distinction between input and output privacy [1], where input privacy is what could be guaranteed by a trusted third party and is generally easier to explain, and output privacy concerns statistical disclosure control and is harder to make precise. Different media, such as video explainers or games that illustrate a toy protocol, can also help, since people have different optimal learning methods. Good real-world analogies also work, but the potential pitfall is that people may take the analogies too far. Another pitfall is that any concept or explainer of good security can also be misappropriated to explain bad security, making it more difficult for users to distinguish between the two.

References

1 United Nations, *The United Nations Guide on Privacy-Enhancing Technologies for Official Statistics*, `https://unstats.un.org/bigdata/task-teams/privacy/guide/2023_UN%20PET%20Guide.pdf`, 2023.

4.4 What are Properties of PET Projects of Social Impact?

Below is a summary of the discussion about properties of projects that have had social impact. Knowing such properties can be useful for building impactful applications in the future.

The discussion identified several notable PET success stories, ranging from end-to-end encryption in messaging apps to differential privacy implementations. The most successful cases share key characteristics: they provide seamless user experiences with zero additional effort required, achieve broad adoption through network effects, and address concrete adversary models. Examples like Signal's encrypted messaging, TOR for censorship avoidance, ad blockers, and Apple's commercial differential privacy rollout demonstrate how privacy technologies can gain critical mass when they either don't increase user difficulty or enable previously impossible functionality. The discussion differentiated privacy-focused technologies from security-focused ones such as HTTPS, Let's Encrypt, and certificate transparency, which were also were noted as successes that improved the overall privacy landscape through widespread deployment.

The group also examined what distinguishes successful from unsuccessful privacy technologies, highlighting that complexity and user friction are major barriers to adoption. Failed or limited examples like Mastodon's complicated setup, encrypted email's usability challenges, and electronic voting systems illustrate how even technically sound solutions can struggle without considering user experience and deployment incentives. The discussion raised important questions about whether imperfect privacy solutions are preferable to no solutions at all, noting the risk of "privacy-washing" for compliance purposes and the challenge that once a technology is deployed, better alternatives face significant adoption hurdles. Corporate interest, financial backing, peer pressure effects, and persistence over time emerged as critical factors, alongside the recognition that some privacy technologies may require accepting trade-offs between privacy, functionality, and adoption potential. The discussion concluded by considering the questions raised by the "Heilmeier catechism," used by the US DARPA agency as a means of assessing the potential of a research project, and how these questions are applicable to developing impactful research.

4.5 Who Should be Financing Public-Good Implementations (and Deployments)? Who is Going to Build, Run, and Maintain PETs Software?

This discussion session combined two questions and we start by addressing the question of who should be financing public-good implementations (& deployments) of PETs.

Before we can proceed, we need to ask ourselves what we exactly mean by "public good." By public good, we mean:

- there is a public benefit (as opposed to a situation where only a single individual, a small group of individuals, or ca ompany benefits from it);
- it addresses a societal problem;
- the PET enables a greater public good, and *creates new value*, and the latter is the main driver for adoption (not just "improved privacy").

With respect to public-good applications, we should view PET software or infrastructure as a public utility, like water and electricity, thus part of the (government-funded) basic infrastructure. We then tried to identify and list possible sources of funding for PET implementations/deployments:

- government (providing continued/sustainable public infrastructure funding);
- (industry) associations (;e.g., Stifterverband in Germany, or G-BA);
- specially created legal vehicles for collaboration (e.g., TMNL in the Dutch banking sector);
- foundations (e.g., SINE);
- (social impact) investors;
- government / EU as a (launching) customer;
- government / EU in the form of research funding, at various TRLs; (continuation is often a problem – a project typically goes to the "PET graveyard" after funding stops);
- crowd-funding / donations;
- multi-national collaborations (e.g., EU-US cyber intel sharing, Eurostat, United Nations).

We also identified several factors that can stimulate the adoption of PETs:

- "spreading the word" about the existence of PETs;
- education about the benefits of PETs;
- raising awareness, for example, through "privacy labels" (like food labels);
- regulation (like GDPR) and its enforcement;
- value creation;
- pressure from society (example: COVID contact tracing);
- cyber threats (international/cross-border).

To answer the question of who is going to build, run, and maintain PET software, we established a list of potential entities:

- foundations (like SINE);
- SME companies (like Cybernetica, Roseman Labs)
- large corporations (like Apple, Google, NTT, Alibaba);
- joint legal vehicles created by stakeholders for collaboration;
- cooperatives;
- public sector (e.g., Data Protection Agencies);
- universities;
- partly government-owned institutes (like Fraunhofer [DE], Alexandra Institute [DK], TNO [NL])
- associations (like WFA for AdTech).

4.6 When Do PETs Work and When Do They Not Work? When Can PETs be Privacy Washing Instead?

Good solutions are not easy to understand, and easy solutions are not good.

4.6.1 When Do PETs Work and When Do They Not Work?

MPC tends to work if you need several partners in your computation. In this scenario, the parties avoid having a single central database that already has all data, but are able to answer queries as if there were one. MPC is useful if you have multiple data holders and everyone agrees that data should not be kept in a single super-database, but there is also a need to compute on that data. However, the parties may not be willing (or allowed) to even secret share their data. Whether MPC would still work if a government changes and the new government is not privacy-friendly, depends on the setup. The number of parties is relevant (for instance, n in the range of 2–5, $n > 10$, contact tracing). In federated learning: the more parties, the better.

MPC may not be applicable in the supplier and business partner scenarios. Compute at one party, verify at n parties is challenging to realise in MPC; in this case, other tech might be more suitable.

There is a difference in tailor-made MPC solutions and general-purpose MPC solutions. Tailor-made MPC is time consuming to implement, can be error prone. However, it can also outperform general-purpose MPC.

In MPC, there can be cases where the participating parties are not equal; e.g., one has data, the other has computational power. Then the possibility of using MPC depends on the trust model: do the input parties trust the computing parties and the whole process?

MPC does not help with other security issues (e.g., input authenticity/integrity). In fact, it might hinder data integrity checks if special measures are not used.

If the government has seen all the data anyway (e.g., in the case of China currently), what is their motivation to utilise MPC or other PETs?

In some societies and countries there is the mentality of "I have nothing to hide" and they expect that this applies to everyone. However, you are not going to reach some of the target groups if you do not use PETs.

Personal data protection and the industry setting have different incentives (push for innovation). There would probably be value in banks sharing their data as even their aggregate information might be compromised. Hence, banks are cautious and might be willing to try out different PETs.

In federated learning one can see more tangible improvement: it is possible to learn on confidential data that was not accessible before. However, classical federated learning is not a PET as it does not protect privacy. If aggregation is secure, it might be considered a PET, but models still leak information. Encrypting gradients does not provide adequate protection, but the combination of homomorphic encryption and differential privacy might be successful.

4.6.2 When Can PETs Become Privacy Washing and the Results be Misused?

Federated learning has a privacy washing feel as you still take all the data in.

MPC for some functionality can also be privacy washing, as it depends what the end result will be used for and what it contains. Also MPC where the users have no control and the computing party decides everything (e.g., what to compute, what to publish). The privacy claim only holds when the computed functionality is restricted and the computing parties obey the boundaries of the allowed computations and the amount of leakage.

Differential privacy without disclosing the parameter epsilon or when epsilon is too big is definitely privacy washing. The central question is whether epsilon is used responsibly.

The elephant in the room is anonymisation to get out of the restrictions of GDPR. Often the methods used are not sufficient, the data is re-identifiable, and data donors' privacy is compromised.

Often synthetic data is generated simply by shuffling IDs. This does not work and can be reversed, but gives organisations the ability to say that they are using synthetic data. This is definitely privacy washing.

4.7 What Would a Comparative Evaluation Framework Look Like for MPC?

Here is the summary from the discussion towards establishing a comparative evaluation framework. The rationale behind this topic was that multiple research areas have a standard set of benchmarks (e.g., datasets or test cases) for evaluation of new developments in that area and a more systematic way of evaluating MPC constructions or tools would be valuable to have.

First, we identified multiple issues we commonly encountered in current protocol evaluations: Parameters, such as network settings, computing power, single vs. multithreaded implementations, etc., often lack clarity. Additionally, there is often ambiguity regarding what is actually measured, e.g., whether runtimes include the offline phase. We shared experiences on papers whose artifacts either contradicted details mentioned in the paper or where significant parts of the evaluation were missing.

As MPC is a very versatile field and covers many different application settings, strict rules that enforce benchmarks in consistent settings will not be adopted. Therefore, guidelines are needed. Authors should aim to adhere closely to the guidelines, and explicitly state and justify any deviations. This is inspired by successful applications of such guidelines in other fields, for example the NeurIPS reproducibility checklist or "datasheets for datasets."

Encouraging authors to be more explicit was a general theme, especially about input parameters and the nature of the actual computations as discussed above. The inclusion of metrics beyond the runtime, such as the number of communication rounds, the total data sent during computation, and, if applicable, solution quality (e.g., accuracy in PPML) should be encouraged. To facilitate this process, providing a common LaTeX template covering essential details and a checklist could be beneficial.

Note that this mostly focuses on the reporting of evaluation parameters, and not on the selection of the evaluation parameters. This emphasis stems from the necessity of establishing reporting standards first. However, there was also support for desk-rejecting MPC papers that evaluate on a single machine without artificial network restrictions.

Regarding guidelines for artifacts, there was a brief discussion on whether just having "high-level code" suffices, or whether MPC engines should also be included. The latter facilitates re-evaluations but might not always be possible, e.g., in the case of proprietary implementations like Sharemind. Nonetheless, publishing "high-level code" should already contribute to reproducibility and makes catching creative benchmarks easier.

Additionally, artifacts should include runtime logs with metadata, such as the network setting (and other parameters), for each data point. This information is crucial to maintain usability when archiving code and data (as mandated in Sweden), and also is useful to independently find answers to questions not discussed in the paper, or to clarify on online and offline benchmark times.

For some sub-fields of MPC, there might be more concrete benchmarking guidelines. For example, PPML papers should always report the accuracy, which needs to be measured on some standard datasets. Private Set Intersection also might allow for more standardized symmetric and asymmetric set sizes. Furthermore, the SPHERE project might be used to both provide computation platforms, and to make existing artifacts accessible. (It did not come up during the discussion, but Dagstuhl also has an artifact publishing platform called DARTS.)

Finally, we also discussed the guideline creation process. An organizer will need to coordinate an initial draft that undergoes iterative validation by lots of authors and reviewers before publication as a paper. Periodic revision of the guidelines every few years will probably be necessary. To promote the use of the guidelines, reviewers could initially encourage authors to implement the guidelines at venues with revision option, such as PETs.

There also was a brief discussion on the chance of pilot projects currently buried on the PET graveyard coming back from the dead when the time is right. For example, this was seen in both ML and MPC priorly. However, it is questionable whether protocols today would still be applicable in the future.

Participants

- Mark Abspoel
Roseman Labs – Utrecht, NL
- Hendrik Ballhausen
LMU – München, DE
- Carsten Baum
Technical University of Denmark – Lyngby, DK
- Marina Blanton
University at Buffalo – SUNY, US
- Dan Bogdanov
Cybernetica AS – Tartu, EE
- Niek Bouman
Roseman Labs – Utrecht, NL
- Andreas Brüggemann
TU Darmstadt, DE
- Kevin Butler
University of Florida – Gainesville, US
- Kasra EdalatNejad
TU Darmstadt, DE
- Vincent Ehrmanntraut
RWTH Aachen, DE
- Simone Fischer-Hübner
Karlstad University, SE
- Mariana Gama
KU Leuven, BE
- Jonathan Heiß
TU Berlin, DE
- Meiko Jensen
Karlstad University, SE
- Liina Kamm
Cybernetica AS – Tartu, EE
- Marcel Keller
CSIRO – Eveleigh, AU
- Ryo Kikuchi
NTT – Tokyo, JP
- Ágnes Kiss
SINE Foundation – Berlin, DE
- John Liagouris
Boston University, US
- Wouter Lueks
CISPA – Saarbrücken, DE
- Kajetan Maliszewski
Technische Universität Berlin, DE
- Kazue Sako
Waseda University – Tokyo, JP
- Sinem Sav
Bilkent University – Ankara, TR
- Thomas Schneider
TU Darmstadt, DE
- Boya Wang
EPFL – Lausanne, CH
- Susanne Wetzel
Stevens Institute of Technology – Hoboken, US
- Rebecca Wright
Barnard College, Columbia University – New York, US
- Hossein Yalame
Robert Bosch GmbH – Renningen, DE