Report from Dagstuhl Seminar 13062

Decentralized Systems for Privacy Preservation

Edited by

Sonja Buchegger[1], Jon Crowcroft[2], Balachander Krishnamurthy[3], and Thorsten Strufe[4]

1 **KTH Royal Institute of Technology - Stockholm, SE,** buc@csc.kth.se
2 **University of Cambridge, GB,** Jon.Crowcroft@cl.cam.ac.uk
3 **AT&T Labs–Research – Florham Park, US,** bala@research.att.com
4 **TU Darmstadt, DE,** strufe@cs.tu-darmstadt.de

Abstract

This report documents the program and the outcomes of Dagstuhl Seminar 13062 "Decentralized Systems for Privacy Preservation". In recent years, a number of concerns have risen about the existence of large, organizationally centralized online services (cloud services, online social networks, repositories, etc). The concerns include risks to users' data from organizational failures and threats to user privacy. In this seminar, the organizers brought together a somewhat more diverse collection of theoreticians and practitioners from industry and academia including social scientists and economists. In keeping with the nature of the interdisciplinary attendees, the organizers also attempted a seminar organization structure intended to promote innovative, cross-discipline working. The results were mixed: some clear agenda setting outputs emerged with some less clear ones.

Seminar 03.–08. February, 2013 – www.dagstuhl.de/13062
1998 ACM Subject Classification K.4.1 Public Policy Issues: Privacy, K.6.4 System Management: Centralization/decentralization, D.4.6 Security and Protection.
Keywords and phrases Privacy, Decentralized Systems, Economics, Usability, Mobility
Digital Object Identifier 10.4230/DagRep.3.2.22

1 Executive Summary

Sonja Buchegger
Jon Crowcroft
Balachander Krishnamurthy
Thorsten Strufe

Distributed and decentralized systems offer more potential resilience to various failures, and, on paper, higher aggregate availability than centralized systems. Centralized management repositories lead to potential risks to users' privacy and the temptation to monetize processing of large aggregates of such data, as seen in systems such as webmail, search and online social networks. Recent years have seen the emergence of projects building prototypes with varying levels of decentralization to reduce these risks. Such systems have not seen great success in contrast to large cloud services. This seminar bought together diverse groups to tackle a series of questions to attempt to answer what may be the root causes of the logjam preventing success of these alternative approaches. There appears to be some consensus amongst at least some groups that there are good reasons for these alternatives. We present here the

Decentralized Systems for Privacy Preservation, *Dagstuhl Reports*, Vol. 3, Issue 2, pp. 22–44
Editors: Sonja Buchegger, Jon Crowcroft, Balachander Krishnamurthy, and Thorsten Strufe

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

output of our group working sessions on these questions. We also provide the reasoning and outcomes of the discussions along with an evaluation of the effectiveness of our mode of working in this seminar.

2 Table of Contents

3 Seminar Plan

3.1 Subject Introduction

Centralized collections of user data have threatened privacy due to data mining and intentional or accidental data leakage to third parties. Online social networks and social media sites are prominent examples, as they attract the lion's share of the Internet users' time today. These recent, Web-based services frequently provide comprehensive personalization, aiming at a precise identification of the individuals using them. While offering valuable services to the individuals on the Web, they collect large amounts of information about the users, including the content willingly uploaded by the users themselves, but more importantly patterns in their preferences and behavior as well as relations to others. All this personally identifiable information is concentrated at a small set of companies that are logically centralized service providers. Large collections of extensive, detailed personal information are needed by these providers, since their exploitation, primarily for targeted advertisement, represents their main business model.

Numerous attempts have been undertaken to counter threats regarding the centralized collection of information. One promising approach is to create services such as online social networks, private data storage and backup, or anonymous content dissemination in a distributed fashion, thus removing the centralized provider with all its knowledge and power. Typically, the gatekeeper functionality of the centralized service provider is replaced by using cryptographic means for access control, metadata-minimizing system design, and other privacy-enhancing technologies to prevent unauthorized data leakages.

While there have been advances on the technical front for decentralized social networks, usability and user acceptance are a challenge. Economics remains a key issue to address head on for any decentralized approach to work. A decentralized approach to privacy-preserving systems inherently means a paradigm shift from today's, mostly Web-based, services. This shift opens a range of research questions in terms of Computer Science (feasibility, scalability, security, new privacy challenges, robustness, resource allocation, resource heterogeneity, efficiency, mobility, etc.) and other disciplines such as Economics, Law, Policy Research, and Sociology. Considering the vast acceptance and ubiquity of these services and their impact on the daily life of individuals, decentralization for privacy is not limited to academic research but needs contributions from other parts of society, such as industry, activists, communities, and policy makers.

We see a number of challenges to be overcome when pursuing the idea of decentralization as a means to increasing the control and privacy of the users. Serving as nuclei for discussions at the seminar, we divided them into three grand challenges User Challenge, Economic Challenge, and Technological Challenge.

3.2 Position Papers

Before the seminar, we sent out a call for one-page position papers to all the participants. The call included the following prompt for the challenges identified for privacy-preserving decentralized systems.

The users. Acceptance and adoption of decentralized services requires a critical mass of users who expect a constant flow of recent, rich content, which conflicts with privacy preservation. How can awareness be raised and to what extent will users actually care about their privacy is one challenge.

Economics of decentralized systems. Centralized systems are highly profitable and monetizing personal data subsidizes such services. Sustainable decentralized systems without monetization potential remains an unresolved challenge. Privacy-preserving queries or introducing differential privacy may help but there are no functional systems yet.

Technological feasibility. Systems requirements like availability, scalability, and robustness or mobility have been discussed but the security of such designs is not well understood. How can functional extensions, such as recommendation schemes, be implemented without access to the entirety of user data and behavior?

The position papers were collected and made available to all participants, with a request to take them into account when introducing themselves and their research briefly at the beginning of the seminar. We asked the participants to state which position paper they resonated, agreed, and disagreed with the most. Additionally, the position papers served to set the scene before the seminar and to extend the challenges and questions already present in the seminar proposal.

With input from the position papers written by the invited participants, the three aspects were further detailed, thus extending the User Challenge to *Users and Usability*, and the Economic Challenge to additionally address *social bootstrapping* (or: roll-out) and incentives.

3.3 Work Sessions

The seminar was organized in sessions of team work, rather than a sequence of presentations with subsequent discussion. Relevant questions for the seminar were identified in an initial expert session and then used for dividing the participants into groups on the three aspects. The seminar addressed each of the aspects separately, discussing economical aspects and users/usability aspects in one separate session, each, and the technological aspects in two sessions.

Mixed sets of experts and other participants were grouped by interest to discuss one of the identified questions from the respective topic and to report their results in a plenary session. The experts on each aspect finally collected these outcomes and established conclusions for their field from the seminar, which are reported in the following sections.

The reasoning behind the division into two types of sessions, called expert and mixed sessions, was to strike a balance between interdisciplinary exchange and deeper discussions between those of similar expertise. This allowed the participants to both learn from perspectives from other fields and to discuss without having to explain a lot of background. The two expert sessions bracketed the set of mixed sessions. The first session of experts had the task to come up with questions to be discussed in the subsequent mixed sessions. The second expert session took place after the mixed sessions had come up with some answers to the posed questions. The task of this expert session was to synthesize what had come out of the mixed sessions and draw conclusions. The mixed sessions in between had participants who joined the groups according to their interest and discussed specific questions.

3.4 Reflections on the Seminar Setup

The structure of the seminar deviated from the canonical Dagstuhl seminar in that there were very few talks, a 5-min no-slide introduction by everyone and few individual presentations by participants. The program instead mostly consisted of interactive group sessions with

discussions on specific topics followed by short presentations of the outcome of these sessions in the plenary. The goal was to encourage knowledge exchange, engaging and coming up with ideas and questions for new research.

While this was certainly more interactive and resulted in some new ideas, it might, on reflection, have been excessive in the emphasis on interaction. What we observed and was mentioned in some of the comments in the seminar evaluation survey pointed to people getting saturated with interaction all day almost daily. It might have been good to switch the mode to presentations even if not as frequently as at a typical Dagstuhl seminar. One effect of the interactive day seemed to be that, at least in some cases, energy that would have been left for informal research discussions over dinner or after, had already been channeled into the work sessions.

To kill two birds with one stone, a set of (reasonably short) tutorials by domain experts for example in economics or other non-Computer-Science fields could provide more background information for group discussions to start at a more advanced level while allowing for some less interactive time.

The position papers elicited before the start of the seminar proved to be very helpful to get an overview of the participants' interests and research. They were appreciated, both by the organizers to distill questions and by the participants themselves to engage with each others' work and perspectives.

Grouping participants by expertise and interest, both self-organized with some slight load-balancing by the organizers when warranted, worked for the most part yet also failed to break up some cliques of people that tended to choose the same groups in several sessions. One clique discussed items not part of the agenda reducing their overall contribution to the seminar. To avoid this, group member allocation might be more strongly enforced by the organizers when needed.

4 Challenge: Users and Usability

Topics expected from proposal: *Acceptance, adoption, and usability*

Additional topics raised in position papers: *transparency, trustworthiness, understandability and take-it-or-leave it EULAs with unnoticed changes, tracking of data disclosure*

4.1 Background

Centralized Social Networking Services enable communication between a group of users, sharing and browsing of content. Since they can collect data about user's actions, they can generate detailed digital dossiers and aggregated profiles of their users, which can be used for targeted advertisement.

The survival and growth of online social networks requires a minimum amount of appealing and fresh content to insure their attractiveness and to guaranteeing frequent return visits of users. They rely on the contribution of participants, publishing their details, opinions, and other user-generated content, and on the visibility of this content to others. Active sharing is commonly limited to a comparably small number of active contributors, whereas the lion's share are primarily passive users who only occasionally, motivated by other posts, share some

content, or comment on discussions. It is essential for the services to reach a critical mass of contributors, and contributions visible to the common user to remain attractive.

User acceptance consequently partly depends on a wide adoption, winning enough active contributors of attractive content, and partly on the availability, and visibility of their contributions. This conflicts with the objective to protect the participants' privacy and to restrict sharing of personally identifiable information to a limited group of actually trusted individuals, only. With smaller user bases, better privacy protection, and lower volumes of attractive content for passive users, a challenge for decentralized, privacy preserving approaches is going to be to attract sufficiently large user bases and available content.

Decentralized approaches may have further shortcomings as compared to commercial services, which exacerbate this situation. Lacking funds to employ professional manpower, they may be characterized by lower usability, fewer extensions, such as games and social apps, and infrequent maintenance to improve their appearance, and absence of marketing campaigns.

Additional protection of each individual's privacy may increase the appeal of such approaches. The question to which extent the average participant of an online social network actually cares about their privacy, and hence, how much the protection of privacy would actually make up for a less content and lower usability remains open.

Beyond the protection from observation by a provider, eliminating centralized storage and control additionally fosters freedom of speech. This freedom, however, raises further challenges, since a lack of control not only prevents censorship, but also deprives the system of the possibility to deal with unwanted, or illegitimate content.

These challenges have to be further explored, to evaluate their impact and aim at potentially coming up with ideas for meeting them.

4.2 Discussion

The expert group on "Users and Usability" identified four questions for further discussions.

4.2.1 Question 1: How to reduce the gap between the complexity of the system and users' mental models?

The intent behind this question is to examine the mismatch between what users expect a system to do and what the system actually does. A good example is persistence of user data: drawing on from real life, users sometimes expect a conversation thread to be ephemeral, but are surprised to find out later that the data related to the conversation is long-lived. Many systems provide fine-grained access control for user data, but the level of expertise needed to properly configure and manage such access control is beyond most ordinary users.

The mixed group discussion recognized that the problem is exacerbated by the fact that different users have different levels of expertise, and a given user's expertise and hence their mental models have evolved over time. This led to the insight that the first step in solving the problem is to develop a reliable way to measure the gap between a user's mental model and the actual system functionality. This may be done implicitly by drawing inferences from user actions, or explicitly, by posing questions to the user. Both of these approaches are not straight-forward and would need to be designed carefully in order to pass the user-acceptance test. The end goal would be to have the system provide guidance to the user in an adaptive manner, depending on the user's expertise and what he/she expects the system to do.

4.2.2 Question 2: How to reduce of the risks of the user's data being misappropriated or used unintended?

The second question was to address the problem of user data being used in ways the user did not intend. This can happen unintentionally, as was the case when a private photo of a Zuckerberg family gathering intended to be shared just with friends became visible to a journalist because of the way Facebook access control works (a photo is visible to the friends of anyone who is tagged in addition to those intended by the sharer). The journalist assumed the photo to be public and shared it further. This example illustrates the consequences of the gap between the (journalist's) mental model and actual system functionality (Facebook's access control for photos with people tags).

Misappropriation can also happen intentionally, when a data collector who goes out of business decides to sell the data to third parties.

The mixed group discussion distinguished between two types of misappropriation.

The first, unintentional misappropriation by "honest but clueless" users can be addressed by improving usability of the sharing process. First, better tools that visualize the extent of sharing to users would alert them if they were about to share with a larger audience than they intend to. This is closely related to Question 1: the type of visualization needed will certainly depend on the user's expertise and mental model, and therefore needs to be adaptive. Second, better techniques for easily selecting the audience of a sharing action can help, too. For example, the data (who is in it) and metadata (where was it taken, who was nearby when it was taken) in a photo can help infer potential sharing targets.

The second, intentional misappropriation by malicious users or data collectors will require stronger protection mechanisms. Currently, the only mechanisms widely in use are legal and regulatory mechanisms which seek to ensure that data collected for one purpose cannot be used for another. There are two other potential approaches one can envisage. The first is the use of trusted hardware. Trusted Platform Modules are widely available for personal computers and server platforms. Hardware security mechanisms for smartphone, like ARM TrustZone, are even more widely deployed and in fact used to ensure features like secure boot. We can build on these features to allow remote platform attestation and configuration verification. Thus, before handing out sensitive data to a remote system, software running on behalf of a user can verify that the hardware and software configuration on the remote system is trustworthy and will not allow the misuse of data. Similarly, advanced cryptographic protocols like private information retrieval are now practical enough to be implemented even on smartphone platforms. Again, neither the use of trusted hardware nor systems that incorporate advanced cryptography is straight-forward. For example, both may limit the ability to run data analytics on user data which is currently used both for monetization (e.g., advertisements shown by Facebook) and for performance improvements (e.g., system health monitoring in systems like Tor).

4.2.3 Question 3: How can decentralized systems reduce the impact of power imbalance while not compromising usability?

Since the issue of decentralization did not crop up in the formulation of the other questions, the expert group discussed what effect decentralization has from the point of view of users and usability. They concluded that the primary difference that concerns the user is the imbalance in power which is manifest in all existing centralized systems: an online social network could arbitrarily change its privacy policies or usage of user data and the users have

no means of protest or debate other than by leaving the system. The network effects make even the possibility of leaving the system less of a free choice (by leaving the system, users run the risk of disrupting social contacts with people whom they care about as long as they remain in the system). This was the rationale for Question 3.

The mixed group concluded that decentralization per se is neither necessary nor sufficient to redress the power imbalance. The pre-requisites for reaching parity of power is (a) standardization of interfaces and (b) governance of the system by a neutral body. They did conclude that a decentralized system is more likely than a centralized system to reach power balance.

4.2.4 Question 4: How to mitigate collateral damage?

The last question was motivated by the fact that social networks significantly extend the speed and scope of rumors. The mixed group was asked to think of ways of limiting the damage from such rumors.

The mixed group divided the problem into three aspects. The first is detection. A system that has the ability to detect that information about a user is spreading can more effectively respond in case the information is (unfounded) rumor. The second is limiting the damage. There are several possible approaches: rate-limiting the spread of information, providing anti-rumor mechanisms like requiring quora and moderation by trusted third parties or representatives of peers, and educating honest-but-clueless users about the impact of spreading rumors. The last point relates to the issue of visualizing the effects of an action as in Question 2. The third aspect is post-damage control. The ability to trace the provenance of data by having an audit trail will help the victim identify how a rumor was spread and respond effectively. Having the possibility of rebuttal by the victim will also help.

4.2.5 Conclusion

Overall, the discussions identified two novel research issues. The first, which arose in Question 1 (and is relevant to Questions 2 and 4), is "How to reliably measure the gap between a user's mental model and system functionality so that the guidance provided by the system can adapt accordingly?". This is an open and difficult research question.

The second is the entire area of Question 4: "How can users mitigate the damage arising from rumors in social networks about themselves?", which opens up many interesting questions as to how to design anti-rumor and rebuttal mechanisms and the pros and cons of throttling the spread of information.

5 Challenge: Economics

Topics expected from proposal: *Sustainable operation, privacy preserving analytics, quality of additional services (recommenders), successful roll-out and incentives*

Additional topics raised in position papers: *Value of commercialization (targeted ads), value of detailed PII, value of privacy*

5.1 Background

Current centralized online social networks are highly valuable businesses. Similar to many web services, they have turned around the model of how charging is done. It has historically evolved from static, embedded advertisement, through search engine's effort to increase click-through rates by profiling and personalization, to the accepted paradigm shift from "The user as a customer", to "The user is the product". These business models are termed "two-sided" markets, since the intermediary provides a service to customers on one side, usually for free in terms of money (but does impact their resources in other way, including eyeballcongnition tie, and network and screen and possibly battery life on mobile devices), and charges advertisers on the other side, money for delivering targeted commercial information to the customer side.

Centralized approaches are based on quite simple economic foundations: The revenue stream flows from advertiser to content/service providers, and the reasons for its success are the two possibilities the digital media offers over traditional mass media.

1. The service provider may know when someone acts on an advertisement (click through) and is capable to collect and analyze quite detailed information about this individual.
2. The service provider knows the demographics of its user base, and hence has the opportunity to perform detailed market research.

Such an environment creates a power imbalance (much like early online banking did in security assurance). The customers get a service for free in return for risking their privacy. Because the service provider wants to do a land grab on all services (healthcare, online shopping, travel, as well as education), they are incentivized to gather more and more details about their users.

This paradigm has evolved from conventional loyalty card services. Individuals in the past had a loyalty card for each different service. Health records were separate from banking. Work records and income were separate from insurance information. Unifying and centralizing all this information introduces serious risks. They introduce an immense attack surface in terms of technical, human and economic weakness, as well as in terms of smallness of the gene pool.

Depriving the systems of their economic bases, by removing the possibility to profile and target advertisement, may introduce economic pressure on decentralized approaches. Economic pressure, while certainly an issue, may not be a fatal obstacle to the provision of the service, as the success of free and open source software has shown in the past. The lack of funding as an incentive, however, may pose to be critical for the allocation of resources, cooperation, and even the provision of content.

Furthermore, decentralizing and encrypting PII may help privacy and reintroduce consumers' rights. However, they may come at the cost of losing oversight of the complete data, thus removing the control to index and locate resources, and to fight unwanted content as well.

Some of these problems may be resolvable by providing privacy preserving queries or trying to apply differential privacy. While numerous efforts have been undertaken, no one has been able to build such a system yet.

Some known methods may help to overcome these challenges, like game theory and mechanism design, experience from networking sciences and distributed systems as well. These issues require discussion, and approaches as well as road maps to aid solving them have to be identified.

5.2 Discussion

The expert group on Economical Aspects identified two complex fields of questions, which then were addressed in parallel by two mixed groups.

5.2.1 Question 1: Are personal data markets viable?

The first set of questions explores personal data markets. The attempt was to try and estimate the value service providers can extract from personal data in the current (ad-sponsored) model with the intention to bound the amount of compensation needed if privacy-friendly systems deprive service providers of this source of income. Likewise, it was attempted to estimate the willingness of users to pay to protect their personal data. Comparing these estimates should yield, ceteris paribus, if a market clearing could be expected or not. Although some price information may be observable in practice, externalities and context-dependence make it hard to interpret these indicators as reliable proxies for the value of personal data.

Pondering the viability of data markets, both groups discussed the actual value of personal data, when exploited for behavioral advertisement. Targeted ads currently make up for a small share of the overall advertisement market only. Even the large players like Google and Facebook can earn only two to three US dollars for behavioral advertisement per user and year. This, however, would potentially already match the cost of a decentralized social networking system. Then again, it certainly is too low to create and maintain a data market between users and advertisers. It is much lower than the revenues cell phone providers and manufacturers can realize. Taking a closer look, the groups detailed that (a) demographic information literally has only negligible value (given the prices charged by commercial services like Spokeo, Rapleaf, and Equifox for such kind of information), (b) information about creditworthiness, income, health, or consumption patterns may be slightly more valuable, and finally (c) the information value decreases over time.

Trying to assess the cost users could be willing to pay for the protection of their data yields a different picture. Protecting their personal information today may be compared to the cost of an insurance that will pay off in future after potential events of data loss. The discussion then led to the conjecture that events of data loss could become so frequent and ordinary that they may not even yield any consequences, nor reputation loss for the culprits any more, as can already be observed at recent examples (companies losing large sets of login/password pairs).

Estimating how much users may be willing to pay for the protection of their data, the groups found instances where users are willing to pay on the order of hundreds of $US per anno for services that promise to protect the online image of a user (ReputationDefender, DeleteMe). This can be supported by the fact that individuals are willing to pay for curtains or to opt out of listings in phone books. However, there is currently not enough data to support claims that enough users would be willing to pay substantial amounts of money for multiple large businesses to form in this area.

A noteworthy observation in the discussion was the fact that users are actually willingly give their data away, if they are convinced it was for a good cause. Individuals participate in focus groups, surveys, loyalty programs (Groupon, Payback, Frequent Traveler programmes) and even disclose their entire browsing history to researchers, if asked. The main concern hence is not the fact that the users lose data, but only for whose benefit and to whom they voluntarily give it.

5.2.2 Question 2: Does a fully decentralized, privacy friendly SNS break even?

Acknowledging this difficulty and recognizing that exploitation of personal data by service providers is a necessary but inefficient way to refinance a centralized infrastructure (and make profits), we came up with a second set of questions that asks if and under which conditions a decentralized infrastructure is viable without advertising. We start with a set of simplifying assumptions, which we relax, one by one, to approach reality. First we look at the steady state (i.e., no transition from or competition with centralized systems) and assume user homogeneity. Most likely, the benefit of the network exceeds its operating cost. The cost may even be small enough to go under the radar of rounding errors in over-investment (this is considered to be controversial). When users are heterogeneous, the cost of some users may exceed their benefit, suggesting that they may exit the network and thereby impose negative externalities on all others by the loss of positive network effects. Such frictions arising from heterogeneity may destabilize the system and so it is crucial to solve the mechanism design problem to align incentives and internalize these externalities. The last step towards reality is to drop the steady-state assumption because not every system that is viable in steady-state might get there if path dependencies lock society into centralized systems. A key problem in establishing a new system it to reach a critical mass, and the most plausible strategy is to leverage existing decentralized systems to share fixed costs and enjoy the network effects of the existing user base. This leads to the question of identifying the right existing infrastructure that could be leveraged for this purpose.

Judging the potential to break even requires understanding of the involved costs and benefits. Both mixed groups identified the core costs to be monetary costs (payments for participation, donations, fees), inconvenience cost (potentially restricted or inferior functionality, viability), development cost, and infrastructure cost (storage, bandwidth, computing). Benefits include increased privacy, psychological effects (satisfaction to participate in or support such a system), infrastructure utilization (availability of data to others, traffic consumed for service), utilization (using the OSN, potential ease and convenience of use), utilization of additional services (recommendations, reputation).

Some of the benefits clearly gain from network effects: beyond a minimum size of the system, each additional user superlinearly increases the benefits (happiness to have built or to have supported the system, utility of using it). Analyzing the system, three different stakeholders with costs and benefits can be identified: (1) the users (who pay monetarily or in degraded service and functional quality and who gain privacy, potential functionality, and the access to resources), (2) the developers and maintainers (who invest time and possibly money to create the system), and (3) infrastructure providers (paying with money or resources, mainly gaining as philantropists or by increasing the number of participants for their own benefit as users). A group made the noteworthy observation that even users who only use the SNS without providing money or any kind of infrastructure resources are actually providing benefit to others, by sharing content, opinions, and votes.

Both mixed groups dismissed the homogeneous case as unrealistic and directly addressed the heterogeneous case of a strong imbalance in utilization (shared and retrieved content vs. shared resources).

Addressing the potential to adopt a decentralized system, a gradual transition, piggybacking on existing systems, seem to be the only possibility. The majority of users is assessed to be unlikely to pay, set up and maintain partial infrastructure, and to migrate to an unpopulated replacement of any existing system. Strong support may either come from subversion (the satisfaction of participating in something subversive) or regulation (legal acts

upon serious incidents when a better solution is readily available).

Several hosts for piggybacking have been identified, and there seem to be three viable ways: (1) leveraging existing offline communities with an interest in privacy (schools, universities, unions) to support the deployment, (2) bridging with existing services (extract data and social graphs from Facebook), and (3) federation by integrating several services over a common interface as an abstraction layer (implicitly choose the most privacy preserving storage substrate through the selected audience).

5.2.3 Conclusion

The experts concluded that data markets could be viable if there were efficient ways for monetization other than advertising. One example would be to switch from users paying with eyeball time to a direct subscription system. Estimates of the cost of a payment system were mad, and it was a relatively modest fraction of current fixed and mobile broadband data access subscription prices. Although seen as certainly viable a discussion on the economic and organizational barriers to bundling cloud and network services concluded that such an approach was problematic in business and economic terms

A key discussion then turned on the need for efficient mechanisms for micro-payment, such as BitCoin. Some such mechanisms are decentralized themselves reducing the risks of merely moving ownership of data from the storage and processing provider to the digital cash mint.

However, considering that data seems not to have much worth placed on it by the users, the viability of such markets remains questionable. It is clear that more research is needed to understand the difference in perception of value of users' content amongst the users, and between users and providers. Studies on risk perception show that it is wildly fluctuating depending on recent positive or negative experiences.

Summarizing the discussions on the chance to break even and retain sustainability, the experts concluded that the choice of the initial set of users is crucial to accelerate acceptance and to achieve information mobility to ease migration. Since the majority of users is expected to be unwilling to pay, a tightening approach is suggested, in which the service could be created at a surplus to achieve a critical mass, and only slowly be burdened by requiring contributions.

It was noted that online services for music and video have moved through several evolutionary stages from piracy, to pay-by-advertisements, to subscription-based on demand systems, which are now highly successful. Indeed, the middle phase did not appear to impact piracy, but the existence of efficient media subscription services appears to cause rapid reduction in content piracy. This suggests that such an economic approach to OSNs, for example, might be equally succesful in reducing abuse: abuse in the sense of loss of privacy of one user by another they don't know perhaps due to incomprehensible privacy settings; and the abuse in the sense of loss of privacy due to OSN operator's monetizing of users' PII.

Privacy preserving analytics were partially covered by a presentation by one current project startup, which, fully centralized, performs analytics on unencrypted data. Privacy is achieved by processing the data under access protection that is guaranteed by a trusted environment (anchored in a TPM), protecting the results to some extent by an approach similar to differential privacy.

There are several places that one can put the users' content in a decentralized architecture. These are not all necessarily victims of the problems discussed in the well-known results that were bought up by several participants, which pointed out the availability and latency

problems of a full p2p approach to decentralization[1]. Various techniques were discussed where payment for storage for cached copies, or payment for encryption for homomorphic crypto-based interest matching for advertisement delivery and for differential privacy of semi-centralized data stores, to offset these availability problems without resorting to centralizing everything again. Such caches are controlled by users' edge devices' (home machine, home router/hub or smart phone as master copy), and could employ multiple cloud service providers, switching dynamically as desired and costed. A full spectrum of solutions in between is feasible, including some non crypted or non-decentralized data, with the corresponding range of business models.

Again, the current situation is not necessarily a good predictor of the future - performance and availability of home user systems in a future with fibre to the home may be very much better than centralized systems one day.

The cognitive overload of paying should also be properly quantitatively compared with the perceived stress of receiving advertisements.

6 Challenge: Technology

Topics expected from proposal: *Which functionality is necessary, at which quality to meet the user requirements? What crypto primitives are available/needed? Feasibility, Scalability, Mobility and Location Privacy, Opportunities.*

Additional topics raised in position papers: *Device limitations/restrictions (lost control over mobile devices, mobile app permission systems), problems of decentralization (observability, traffic analysis, DoS), novel challenges/threats from decentralization (integrity, end-point correlation, availability), metrics to measure privacy, primitives for privacy preserving recommendation, to establish trustworthiness, decentralized trust (evidence-based trust, reputation with privacy)*

6.1 Background

Privacy threats have risen along with the increase of large-scale data collections not only of user-provided content but also of automatically collected personal, relational, and behavioral data, as exemplified by online social networks, credit and loyalty cards, or tracking on the Web. These data collections are concentrated at (logically) centralized service providers, thus intentional and accidental leaks of private data to third parties can have significant impact.

There has been a lot of research activity on privacy over the last decade, especially in the Computer Science community. A main outcome there has been the development of several privacy-enhancing technologies (PETs) such as onion routing, mix-nets, or anonymous credentials. More recently, there have been approaches that propose to break up provider-dependent centralized data collections and return the control over their data to the users themselves by decentralizing systems and replacing the gatekeeper functionality of a centralized provider by technical means such as cryptographically enforced access control and metadata-minimizing system design.

[1] See for example "High availability, scalable storage, dynamic peer networks: pick two" by Blake, Charles and Rodrigues, Rodrigo, in HotOS 2003.

The main focus of these efforts has been on decentralizing social networks. Regular, centralized online social networks are a particularly good example for this problem as they have an extremely large user base and collect information in addition to what the users upload about themselves. They have data on what users say about other users, whom they interact with, and other behavior also on third-party sites thanks to tracking and functionality such as liking content on the web.

Different models for decentralization have been proposed. One common approach is to store user content in a distributed system, such as a peer-to-peer network with replicas or at least a collection of independent servers instead of a single-provider and control. There are design challenges at several levels in terms of security, availability, scalability, robustness, new privacy issues, usability, etc. for any decentralized service or application, underlying storage and network topology, as well as trust relationships. One particular challenge is to compensate for the privacy-preserving and security functionality that does exist in centralized single-provider services that can, for example, hide behavior data from other users.

It is not clear to what extent decentralization can be both feasible and beneficial for keeping the system functionality and preserve privacy. For instance, there are trade-offs in terms of provider independence and resource allocation and trust management for fully decentralized systems versus centralized systems with cryptographically protected user data versus federated solutions.

The potential and the limits of decentralization as a means to enhance privacy have to be explored. After an overview of the state of the art is established, we will be able to detect gaps and thus determine what a research roadmap would be to bridge these gaps. While there will continue to be a need for diversity in research approaches, areas of synergy have to be identified.

6.2 Discussion

The expert group for the Technological Aspects identified six questions, divided into two blocks for separate mixed sessions.

6.2.1 Question 1: What are the important design goals of decentralized privacy-preserving systems?

The goals can be separated into several categories. The first category encompasses privacy requirements. At a high level, the concept of information self-determination, i.e., that users should be in control of their data, seems like a good starting point; however, there are many details that affect the design of the system. For example, it is not clear what level of resistance to traffic analysis is necessary or desirable.

The second category is the utility goals: what are users expected to accomplish through using the system? Existing OSNs provide a variety of different communication functionalities, from sharing simple status updates, to finding friends, to monitoring for malicious content.

The final category concerns the goals of the users of the system. Here, an understanding of what kind of individual goals each user might have is required; for example, some users might prefer to know when their data is shared, some users might prefer not to be tracked, and so on. Important questions include: how do users express these goals to the system? How can a system realize the goals? And how do can those conflicts between the goals of users be resolved? The latter problem touches on the complexity of data ownership in OSNs.

Different systems will have different goals; a crucial task is to create some form of taxonomy of goals to understand their individual properties, as well as the relationship between them. In terms of individual properties, it's important to understand to what extent is an individual goal realizable, deployable, or commercializable. Other concerns might be whether such a goal would lead to novel research and published papers, and more fundamentally, how well would this type of goal align with the type of aspired society. In terms of relationship, it is necessary to identify which goals are mutually incompatible.

Providing a complete taxonomy within a single session was deemed to be impossible. However, different dimensions and approaches could be identified.

The dimensions along which the goals can be measured are

- whether the goal is for the individual or for a collective
- whether the goal supports freedom of action or information flow control
- whether the goal seeks to embed social norms into technology, or rather, change social norms through technology

These dimensions are not necessarily independent, and they may not always be applicable, however, they represent good guidelines.

Two different approaches to identify goals are top-down and bottom-up: Goals can be defined either with powerful adversaries, or idealistic protection objectives in mind. Another, sometimes potentially more useful and pragmatic way is to identify users and stakeholders as well as their actions and needs, to collect specific requirements and derive goals from them. The second approach partially is motivated from the insight that current practice seems to be to design systems and protocols with certain properties and derive the goals later, instead of developing systems towards actual needs. A simple exemplary approach is to define personae and the actions they take and to analyze how they'd be affected by design decisions.

6.2.2 Question 2: What are the threat models that need to be considered?

There are various types of threats to the security and privacy of a decentralized system; given a set of goals and a potential set of threats, attempts can be undertaken to realize a system that satisfies those goals in the presence of an adversary.

To understand the space of threats, some sort of taxonomy is required. In particular, it is necessary to understand what kind of threats we have to worry about – some might be realistic today, some might become significant in the future, and others might be artificial. In addition to well understood threats, there may be research problems to identify in exploring new types of threats that are specific to OSNs.

Defining privacy threats in the scope of social applications is not straightforward, since it is hard to differentiate between the intent of the user to share some information, as opposed to the dangers of others actually retrieving it. Arriving at a taxonomy of threats necessitates the definition of a hierarchy of categories to avoid simply filling a very large matrix of categories. Both ways, though, do not seem feasible within a single discussion session, and only first steps towards this goal are taken.

Three dimensions have been identified for a taxonomy:

1. The system architecture,
2. Stakeholders
3. Assets.

Considering the categories, the type of system has to be defined and basic differences identified (e.g., centralized vs. decentralized). Further categories are the stakeholders, specifically the actors and adversaries, as well as assets that need to be protected (from directly identifying over pseudonymous to location information).

The stakeholders have to be discriminated and defined, and it is necessary to define who can actually be trusted, and how they are incentivized. The choice of adversaries to address has several sides. The honest but curious adversary may be the most realistic and hence should be taken into account at first. Considering the most extreme cases (like nation states, malicious providers, of even organized crime) yields understanding and lessons, and it may make sense to aim for the strongest but usable protection. Cases with weaker adversaries are valuable, too, though, since they may reflect reality better.

To judge the threats and order them by importance for the sake of prioritizing them, their potential harm, their likelihood of being realized, and finally the effort or ease to fix them have to be considered. Privacy Impact Assessment is a methodology that could potentially be applied for this purpose.

6.2.3 Question 3: How do we ensure that a system has a good chance of seeing the light of day?

A clear prerequisite to the adoption of a system is the filling of an unfilled need. It is not necessarily clear that privacy alone is compelling enough as a need to accumulate a large user base, especially when trade-offs are introduced. One potential strategy may be to appeal to a niche community, at least initially, rather than try to be everything for everyone.

Even so, adoption can be slow and has a bootstrapping issue. The important question here is whether we can leverage some existing infrastructure to ease this bootstrapping burden.

Finally, there are many platforms that people use for accessing OSNs, including desktops, mobile phones, thin clients (web browsers). What are the constraints that these platforms create for the design of the system?

Privacy is not in itself an unmet need by common understanding. It may attract a few, but other reasons to change are needed for the majority of users to actually change their service. Factors of scale, like usability, which can only be guaranteed with a large number of highly qualified developers, as well as network effects actually are strong antagonists of change. Achieving the same functionality, usability and reaching the network effects, privacy as a matter of fact can be considered a compelling property. A successful deployment, however, is much more likely with additional, complementary or innovative functionality.

Piggybacking on other systems to kickstart the acceptance is conceivable in four different ways. A new system can be bootstrapped out of an existing or several existing systems. One possible way is to either harvest the content out of existing systems and replicate it in the novel platform to provide a set of interesting content and make the transition easy for the users. This, however, is impossible for institutional approaches, since the licenses of existing systems prohibit it and the providers are very keen on protecting the content they collected from the users from competition. Another approach could be federation: allowing for seamless integration, system could run behind a single interface in parallel, and for each act of sharing, the most secure medium that reaches all destinations could be chosen automatically. Another

possibility is to bootstrap the system from existing infrastructures, like eduroam[2] (with several hundreds of thousands of users), shibboleth[3], or even SIM card infrastructures of mobile network providers. The third possibility is to leverage existing organizations with an inherent interest in communicating at a minimum level of confidentiality. Data about minors or schools is a good example. But organizations with interest in confidentiality and large numbers of members, such as Unions, are likely to act as amplifiers, as well. The fourth opportunity is to leverage fields of applications that already have large user bases, like CSCW, for instance. Rationale behind such attempts needs to be to reach critical number of users as core by such a subversive ways, to spawn and achieve network effects, to make transition more attractive to the large majority of users.

Regulation is an entirely different factor. It seems useful to develop a secure and viable solution as an alternative for the case when more critical incidents of data loss happen, in order to be able to offer it to regulators as an example, if even only to define the properties that can be required by regulators consequently.

The final sub-question raises the issue of future devices and their restrictions: Mobile devices with locked operating systems are quickly gaining market share, and users on PCs and notebooks are decreasingly willing to actually install applications and background daemons. Future deployment of TPM could exacerbate this. This results in the demand to develop purely browser-based systems, and to address loss of control over devices. It doesn't seem sensible to protect user data on the application layer of (D)OSN from players like Google, or Apple, if the protection is implemented on a device that runs Android/iOS and hence is under control of these entities on a lower layer, anyway.

6.2.4 Question 4: What are the building blocks for decentralized privacy systems?

A sub-question is whether decentralization itself is an essential building block for privacy, given the research into privacy-preserving protocols on top of, for example, cloud computing. For building blocks that we identify, we need to understand the costs and benefits of these tools, in terms of privacy guarantees, performance, the ability to generate revenue, the underlying feature set, and so on. An important question is how such blocks may be securely composed, since in general, the security of composition of functionalities has been difficult to achieve.

A related question is how can we encourage our immediate community to build reusable building blocks so that we can build on each other's work, rather than starting from ground zero each time. Are there barriers to such sharing that we can identify and perhaps address?

The question was slightly adapted to address the potential architectures of privacy preserving social applications, their respective costs and benefits, and how they could be separated into building blocks that compose well.

Decentralization in this context is defined as distribution of storage and control over several authorities, or providers. Scopes of decentralization range from centralized over hybrid to entirely decentralized systems (Facebook, Diaspora, Peerson or Safebook). Federation may be used in a way to avoid putting all eggs into the same basket, and in this case can represent decentralization, too.

To achieve a better understanding, it is necessary to understand reasons to decentralize or

[2] https://www.eduroam.org/
[3] http://shibboleth.net/

the effect of decentralization first. Stronger control may be a motivation for decentralization, as well as for increased performance or reliability. Taking a closer look it becomes apparent that these points are actually not achieved in currently developed and proposed solutions, and they may actually be quite hard to achieve. Further reasons are to avoid being sued (the difference between Napster and BitTorrent and their greatly different history being the lack of an institution in control of the service), and avoiding a single point of failure with respect to collaboration for censorship, or "lawful interception" as defined by regimes and governments of vastly different natures. A final property of decentralization is that they be future proof: while centralized services can run into financial problems or be sold and hence quite drastically change their licenses, this threat does not affect decentralized systems.

The advantages come with drawbacks, and being both censorship proof and confidential service, the threat of abuse is high, and may actually have disincentivizing effects for large fractions of users.

Defining building blocks is actually more difficult than expected. It is neither clear if the developed systems should be general purpose OSN or if its better to create targeted solutions in order to understand features, concepts, and resulting properties. Nor is it straightforward to decide to which extent the building blocks need to be adaptable, keeping in mind the understandability and usability of the respective API.

Some building blocks have been identified nonetheless. The following list may not be comprehensive but serves as a good starting point. The group enumerated them to be: storage (not necessarily distributed and highly configurable), registration (account creation, identification, key-escrow), profile management (content publication, audience selection spanning authorization and access control, potentially by crypto and key-distribution), secure user discovery, a crypto library (integration of useful primitives, usable and comprehensive), notification mechanisms, connection establishment and NAT traversal, technical bootstrapping. Further useful functionality was identified to contain a voting mechanism, partial message ordering, and mechanisms for anonymous communication, even though it wasn't quite clear how general those were.

A final discussion revolved around the question of how reuse could be encouraged and made possible. The main obstacles to this have been identified to be of psychological nature, but good and well documented code, exact and understandable APIs, possibly even the provision of simulation or numerical models could help, to aid the people integrating building blocks understand and assess the properties of the composed system.

6.2.5 Question 5: How do we evaluate system designs?

Systems can be evaluated from a variety of different metrics and using a variety of approaches, from formal analysis to experiments to field deployments. It may not be productive to identify the space of all possible evaluation metrics and strategies, since many of them are not specific to decentralized privacy systems. Instead, we should focus on identifying what types of evaluation is especially important for decentralized privacy-preserving systems, and in particular, what types of privacy metrics might be relevant. Are there existing ones or do we need to define new ones? It would also be fruitful to identify any challenges and barriers to evaluation that privacy-preserving systems create, such as the difficulty of obtaining realistic data sets while preserving individual privacy.

The initial, motivating observation of this discussion was that no consensus nor accepted approach exists, and many previous and current studies lack rigor, realistic assumptions, and reproducibility. This is exacerbated by the fact that comparing two systems by itself

is complicated, and only convincingly possible if one system is a superset of the other, or one dominates the other entirely. There is no agreement on a set of metrics that measures privacy in a convincing manner, let alone in cross-domain scenarios, and thus evaluation is difficult and the results questionable. Even just cultural differences have been pointed out to be a cause for complication, since they define what is acceptable in different cultures, and hence may have an effect on how the protection of different systems is perceived.

Another major obstacle for evaluations is the lack of good data sets that define assumptions and environment. In an attempt to sketch a solution a possibility could be to send academic interns to companies in possession of the data. The interns then could analyze the data, specify exact datasets and their descriptions, implement and perform analyses and evaluations on site and publish only the results. Other could then reproduce the results knowing the specifications by sending their evaluation code to the companies, who would run it and again provide only the results.

6.2.6 Question 6: What is the state of the field?

We should identify which questions might have already been answered well in the research literature (or in practice), and which questions are deserving of future research. It would be useful, though perhaps controversial, to perform a retrospective analysis of existing research and identify directions that we feel are particularly compelling, as well as directions that seem less scientifically interesting or that are based on hard-to-justify assumptions.

Acknowledging that the entirety of the state of the art was too broad to be fully covered in a single session, the team discussed the general situation of the research fields that have to come together to provide privacy preserving services and their integration. It came to the conclusion that while an impressive body of work on crypto primitives, algorithms, and protocols exist, the lack of understanding for the actual requirements and functions led to only partial applicability. This problem has additionally been identified to affect the systems solutions as well, which have seen a large variety of proposals for technological solutions, yet no notable adoption by the users, so far. This conclusion was encountered with the observation that no solution has ever been successful at first shot, and that it's usually followers, who succeed, not the innovators.

6.2.7 Conclusion

The six questions of the technological challenge turned out to be too broad for conclusive discussions.

Taxonomies for goals, threats, and concepts including their properties are needed and should be provided by and to the community. Initial dimensions for classifications have been sketched and discussed, which seem to be (at least) a good starting point. *Goals* could be classified by the cardinality of their subjects (individual vs. collective), by their data control (freedom of action vs. information flow control), and their social impact (implementing or changing social norms). Potential classification dimensions for *threats* are the system architecture (centralized vs. decentralized, type of access control scheme), capabilities and goals of stakeholders (who are actors, who/how powerful are adversaries), and which assets are to be protected (explicitly/implicitly shared data; directly identifying, pseudonymous, location information). Classifying the functional concepts, finally, will greatly depend on assumptions and requirements, and the taxonomies hence will naturally address a combination of the dimensions above.

It has to be noted that decentralization by itself does not directly lead to increased privacy. Even implementing systems that enforce access control over explicitly shared data, decentralized systems can be highly vulnerable to traffic analysis attacks, partially due to the lack the mixing property of centralized systems.

It is not quite clear if a general purpose social networking service should be the aim of decentralization efforts, or if it is more promising to start with specific, more targeted services. An initial, preliminary list of general building blocks for both cases seems to include functionalities for: registration, profile management, secure user discovery, notification, general purpose storage, comprehensive crypto, bootstrapping, and connection establishment and NAT traversal.

7 General Observations, Questions Raised, Open Problems

Privacy by itself is not an unmet need, and providing privacy in an otherwise identical system will not yield widespread adoption. Privacy can, however, be a compelling property; and offering privacy protection in a system with nearly-identical usability and an attractive set of features most probably can be a strong success factor.

Successful propagation of a novel, privacy preserving, either partially or completely decentralized service will thus depend on additional factors. Reaching a critical mass requires the participation of a significant number of initial users, before network effects and word of mouth will help increase general interest. Leveraging existing infrastructures and social structures could help attract this initial user base. Existing infrastructures may be current services, integrated by federating user interfaces, or existing authentication services of large institutions and their conglomerates. Existing social structures may be institutions with the need to communicate at high levels of privacy protection, like schools, or unions.

General observations regarding the assumptions were that systems need to address the heterogeneity of users with respect to their activity and sharing behavior, as well as their need to persist beyond initial startup phases, in which the resource provision potentially can be hidden in the margins of other services, to a sustainable existence. The latter may only be possible if the systems either allow for monetization (beyond the almost negligible advertisement market), or implement an internal incentive scheme to attract sufficient resource providers.

Recurring themes from the technological perspective were the different promises and drawbacks of trusted platform modules (TPM) and both usage and control restrictions of future devices. TPM promise guarantees of trustworthiness: users may be able to restrict and verify the absence of chances for downstream abuse, which may help deciding which information to share with whom. Analyzing changing habits of the users it can be assumed that soon the vast majority will access services over mobile devices or browsers. Both cases prevent installing client software that would constantly run in the background (and in case of laptops and PCs this starts to be frowned upon) to provide services to others. The mobile devices raise the even more important question of how the users can regain control without having to jailbreak or root them, to avoid running a seemingly trusted system on a platform that is controlled by one of the potential adversaries.

Several novel research questions, which are addressed partially at the computer science communities and partially beyond, have been phrased. Questions for our own community covered HCI, economics, and technological fields.

- Oversharing frequently happens due to mistakes and misunderstood controls. *How can*

we measure the gap between the mental model of the users and the reality? How can we adapt and encounter it, to achieve the ***Principle of least astonishment****?*

- More specifically: *How can the de-facto audience of a post be visualized, how can audience selection be simplified?*
- Upon incidents: *How can users mitigate the damage arising from rumors or libeling content?*
- The community is developing large numbers of systems, algorithms and protocols: *How can re-use be fostered; How can sensible and useful crypto-libraries be provided; Which scope should systems (and hence the shared building blocks) target?*
- Considering the introduction of novel, more secure services: *How can existing social structures and infrastructures be leveraged to bootstrap a novel system?*
- The discussions yielded the insight that the behavioral advertisement market is rather negligible, but within the light of this fact: *Are monetary or other incentive markets viable, and how is the sustained operation of a novel service possible?*

Questions beyond the field of expertise of the participants address mainly legal matters:

- *Considering a fully decentralized system, who is responsible (read: can be sued) for its operation, content, and offences committed using it?*
- *In the global context, who owns the data shared on social networks, who has copyright, right of use, right of deletion?*
- *Is plausible deniability realistic protection for institutions running TOR- or Freenet nodes?*
- *To which extent is integration with existing services (federation, use as storage substrate, retrieval of data stored within) legally acceptable?*
- *Could eduroam/shibboleth be extended to global identification services; at which complexity?*

Participants

- Jonathan Anderson
University of Cambridge, GB
- N. Asokan
University of Helsinki, FI
- Rainer Böhme
Universität Münster, DE
- Nikita Borisov
Univ. of Illinois – Urbana, US
- Sonja Buchegger
KTH – Stockholm, SE
- Ramon Caceres
AT&T Labs Research – Florham Park, US
- Jan Camenisch
IBM Research – Zürich, CH
- Jon Crowcroft
University of Cambridge, GB
- George Danezis
Microsoft Research UK – Cambridge, GB
- Claudia Diaz
K.U. Leuven, BE
- Vijay Erramilli
Telefonica Res., Barcelona, ES
- Simone Fischer-Hübner
Karlstad University, SE
- Paul Francis
MPI for Software Systems – Kaiserslautern, DE
- Ian Goldberg
University of Waterloo, CA
- Artur Hecker
Huawei Technologies, DE
- Urs Hengartner
University of Waterloo, CA
- Jaeyeon Jung
Microsoft Res. – Redmond, US
- Mohamed Ali Kaafar
INRIA Rhône-Alpes, FR
- Gunnar Kreitz
KTH Stockholm, SE
- Balachander Krishnamurthy
AT&T Labs Research – Florham Park, US
- Leonardo A. Martucci
Karlstad University, SE
- Bart Preneel
KU Leuven, BE
- Stefanie Roos
TU Darmstadt, DE
- Krzysztof Rzadca
University of Warsaw, PL
- Hervais-Clemence Simo Fhom
Fraunhofer SIT – Darmstadt, DE
- Thorsten Strufe
TU Darmstadt, DE
- Paul Syverson
NRL – Washington, US
- Claire Vishik
Intel – London, GB
- Marcel Waldvogel
Universität Konstanz, DE