Report from Dagstuhl Seminar 15151

Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations

Edited by

David Hutchison[1], Klara Nahrstedt[2], Marcus Schöller[3], Indra Spiecker gen. Döhmann[4], and Markus Tauber[5]

1 **Lancaster University, GB,** d.hutchison@lancaster.ac.uk
2 **University of Illinois at Urbana-Champaign, US,** klara@illinois.edu
3 **Hochschule Reutlingen, DE,** marcus.schoeller@reutlingen-university.de
4 **Goethe-Universität Frankfurt, DE,** spiecker@jur.uni-frankfurt.de
5 **AIT Austrian Institute of Technology – AT,** markus.tauber@ait.ac.at

Abstract

Dagstuhl Seminar 15151 entitled "Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations" brought together researchers from different disciplines in order to establish a research agenda for making future services in our increasingly connected world more resilient and secure, as well as addressing privacy. The participants came from a range of disciplines covering the techno-legal domain, resilience and systems security, and socio-technical concerns. The use case domains that were discussed during the Seminar covered the Internet of Things (IoT) as well as Cloud-based applications in which flexible service composition is a crucial element. From a starting point covering the "big picture", the legal viewpoint, the technical viewpoint, and the organisational viewpoint, we derived initial research questions in small groups, and the questions and issues arising were then consolidated and refined. The groups discussed the issues in depth and have produced the report and the research agenda contained here.

Seminar April 7–10, 2015 – http://www.dagstuhl.de/15151
1998 ACM Subject Classification C.2 Computer-communication networks, J.4 Social and behavioural sciences, K.4 Computers and society, K.5 Legal aspects of computing
Keywords and phrases Resilience, security, privacy, legal aspects, networked systems, organisations, society
Digital Object Identifier 10.4230/DagRep.5.4.1
Edited in cooperation with Simon Oechsner (NEC, simon.oechsner@neclab.eu)

1 Executive Summary

David Hutchison
Klara Nahrstedt
Marcus Schöller
Indra Spiecker gen. Döhmann
Markus Tauber

License Creative Commons BY 3.0 Unported license
© David Hutchison, Klara Nahrstedt, Marcus Schöller, Indra Spiecker gen. Döhmann, and Markus Tauber

This report documents the programme and the outcomes of Dagstuhl Seminar 15151 on "Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations".

Except where otherwise noted, content of this report is licensed under a Creative Commons BY 3.0 Unported license
Assuring Resilience, Security and Privacy for Flexible Networked Systems and Organisations, *Dagstuhl Reports*, Vol. 5, Issue 4, pp. 1–17
Editors: David Hutchison, Klara Nahrstedt, Marcus Schöller, Indra Spiecker gen. Döhmann, and Markus Tauber
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

The main objective of the Seminar was to bring together researchers from different disciplines in order to establish a research agenda for securing services-to-come in our increasingly connected world. The backgrounds and interests of the participants included i) techno-legal, ii) resilience and systems security, and iii) socio-technical topics. The use case domains that were discussed covered the Internet of Things (IoT) as well as Cloud-based applications in which flexible service composition is paramount. We started the seminar using four introductory talks covering respectively the "big picture", the legal viewpoint, the technical viewpoint, and the organisational viewpoint. From this beginning, we derived initial research questions in small groups, and these questions and issues arising were then consolidated and refined into the resulting material that is presented below.

The opening speakers were the following:

- Helmut Leopold, Head of the Digital Safety and Security Department at the Austrian Institute of Technology, who presented the "big picture", i.e. where our connected world is heading;
- Burkhard Schafer, Professor of Computational Legal Theory at the University of Edinburgh, who presented his viewpoint on legal challenges within our ever interconnected society;
- Thilo Ewald from Microsoft Deutschland GmbH, who explained his viewpoint on the organisational challenges in today's world;
- Marcus Brunner, Head of Standardization in the strategy and innovation department of Swisscom, presented his viewpoint on technological developments in designing and building flexible networked systems.

From this starting point we derived initial research questions in small groups. The organising team reviewed intermediate results and re-balanced groups and most significantly identified the core questions to work on. The groups were between 4 and 6 people at any time, and a good balance was maintained across the representatives of legal, organisational and technological experts and between the groups. The resulting questions and issues were:

1. How to enable Resilience, by design, of composable flexible systems [1]?
2. What is the role of law in supporting resilience, privacy [2] and security?
3. Traceability of (personal and non-personal) data in service provision?
4. How can we improve the perception of assurance [3], privacy, security and resilience for the end-user?
5. What constitutes a security problem?
6. How to deal with unforeseen new context of usage?

These questions were crucial, in that they formed the basis for the bulk of group discussions throughout the second and third days of the Seminar. Therefore, the organisers took great care – and a great deal of time during the first evening – formulating these questions, together with the related issues. At the start of the second day, these questions and issues were presented to the groups, who were invited to comment on them. The groups were invited to add their own interpretation, and to identify additional issues during their discussions. During the subsequent periods – broken up by refreshments and lunch – the organisers checked that the groups appeared to be productive and harmonious (which on both counts they turned out to be). Each group was asked to record the essence of their discussions, and conclusions, and to pass these to the organisers by the end of the Seminar. Every group did some additional work after the Seminar, and the report assembled here reflects the hard work of the participants as well as the organisers, during the Seminar itself and in the days that followed.

References

1 James P. G. Sterbenz, David Hutchison, Egemen K. Çetinkaya, Abdul Jabbar, Justin P. Rohrer, Marcus Schöller, and Paul Smith. Resilience and survivability in communication networks: Strategies, principles, and survey of disciplines. *Comput. Netw.*, 54(8):1245–1265, June 2010.

2 Burkhard Schafer. All changed, changed utterly? *Datenschutz und Datensicherheit – DuD*, 35(9):634–638, 2011.

3 Aleksandar Hudic, Markus Tauber, Thomas Lorunser, Maria Krotsiani, George Spanoudakis, Andreas Mauthe, and Edgar R. Weippl. A multi-layer and multitenant cloud assurance evaluation methodology. In *Cloud Computing Technology and Science (CloudCom), 2014 IEEE 6th International Conference on*, pages 386–393. IEEE, 2014.

2 Table of Contents

3 The report

3.1 How we ran the Seminar

Opening talks took place on the first morning; on the first afternoon we ran pre-selected groups to produce candidate questions, which were consolidated by the organisers during the first evening.

Initial group setup, at least one organizer was part of the groups (where the organiser responsible for each group is identified in square brackets):

1. Brunner Balaban Alshawish Mauthe Tsudik [Tauber]
2. Leopold Raabe Fischer Sterbenz Varga [Hutchison]
3. Ewald Lyles Sorge Stiller Weippl [Spiecker]
4. Pallas Kadobayashi Kirby Smith Schaeffer-Filho [Schoeller]
5. Schafer Bhatti Delsing Bless Dan [Oechsner[1]]

We readjusted the groups for the second day, as follows:

1. Brunner Raabe Alshawish Mauthe [Tauber]
2. Leopold Balaban Fischer Sterbenz Varga [Hutchison]
3. Ewald Sorge Stiller Tsudik [Spiecker]
4. Pallas Kadobayashi Kirby Smith Schaeffer-Filho [Schoeller]
5. Schafer Bhatti Delsing Bless Dan [Oechsner]

There followed in-depth discussion of research questions that the groups themselves chose freely based on the above list (the organisers checked that there was sufficient disparity across groups).

At the end of each session, there was a presentation of discussion outcomes from each group.

3.2 Introductory Talks

The introductory talks and other presentations can be found in the appendix:

- Big Picture: Helmut Leopold
  As head of the department Digital Safety and Security at the Austrian institute of technology Helmut Leopold supervises research agendas in multiple fields. He presented his view on the general directions and major trends in research and society. Those included ICT Trends in the „after-broadband century", the Security problem, the Shift in user behaviour and the IT industry problem.
- Legal Viewpoint: Burkhard Schafer
  Burkard Schafer is Professor of Computational Legal Theory at the University of Edinburg and co-director of the Joseph Bell Centre for Legal Reasoning and Forensic Statistics. As such he operates on the intersection between law, science and computer technology. He presented his view point on gabs in this particular intersection and requirements regarding legal definitions. He also pointed out that security research ethics need to be established to allow research a structured way of publishing their empirical findings.

[1] Simon Oechsner supported the organisation team in the absence of Klara Nahrsted.

- Organisational Viewpoint: Thilo Ewald
  Thilo Ewald is Micorsofts Cloud Delivery Executive for Germany. As such he has an overview of the technologies to come and customers requirements and anticipations. He presented the challenges of deploying scalable global office solutions and how security relates to scalability and connectivity.
- Technology Viewpoint: Marcus Brunner
  Marcus Brunner is responsible at Swisscom's Strategy and Innovation unit for standardisation issues which he is leading. He presented technology and customer expectation issues. Focusing on telco provider issues and concluding that customers see security often as part of the telco provision.

Additional presentations were done on:

- Connecting Legacy Systems
  Jerker Delsing is Professor at Lulea University of Technology. He talked about the problem of connecting legacy technology to the internet of everything and supported this with solutions [1] from the arrowhead project[2].
- Address-space, Routing and Mobility
  Saleem Bhatti is Professor at the University of St Andrews. He talked about problems related addressing in the internet of everything and supported this with solutions [2] from the ilnp initiative[3].

3.3 Research Questions

The research questions produced after the opening presentations, which formed the basis for the core group sessions in the Seminar, are reproduced here:

1. How to enable Resilience, by design, of composable flexible systems (new architectural models)?
   - How to overcome complexity?
   - Responsibility
   - Risk-management and assessment
2. What is the role of law in supporting resilience, privacy and security (technology regulation)?
   - How to enforce law – e.g. automated algorithmic law enforcement?
3. Traceability of (personal and non-personal) data in service provision
   - Anonymization and de-anonymization of data
4. How can we improve the perception of assurance, privacy, security and resilience for the end-user?
   - Can we build an economic model for security and resilience?
   - The role of trust in assuring security
5. What constitutes a security problem (model, define, measure, . . .)?
   - Do we need an ethical framework for handling findings from security research?
6. How to deal with unforeseen new context of usage?
   - Shifting responsibility for data

[2] http://www.arrowhead.eu/
[3] http://ilnp.cs.st-andrews.ac.uk/

Each group chose freely from amongst these questions – with the associated issues – to discuss in depth the topics that they find important and can recommend be included in a research agenda.

A summary of the findings for each research question / topic is given below. This is followed in the Appendix by a presentation of the report that each group produced, providing a correlation of the outcome of the Seminar.

3.3.1 Question 1

How to enable Resilience, by design, of composable flexible systems (new architectural models)?

- How to overcome complexity?
- Responsibility
- Risk-management and assessment

The first question to be discussed was how to specify resilience at the enterprise level and how to map this into the system layers and mechanisms.

The group recognized that this is similar to the Quality of Service (QoS) mapping issue that was a subject of considerable research during the 1990s. Specification of desired qualities has to be done at the enterprise (or application level) in a form or in a language that is understood by the end-user, using for example the so-called Olympic levels of Gold, Silver or Bronze, or alternatively using some QoS classes that implicitly encapsulate the desired QoS properties – such as interactive applications as opposed to file transfer (where in the former case, delay and jitter must be minimized, but in the latter, these are less important). In order to communicate the QoS specification into the network, a mapping has to be made from the high-level statement into the corresponding network parameters or metrics (thus, for example, delay, jitter, packet loss and so on). Research into the specification of resilience and the mapping into appropriate metrics is ongoing (notably by some participants in this Dagstuhl seminar), and builds on research that was carried out in the EU Framework Future Internet ResumeNet project. One of the key issues is what metrics to use at the service and topology levels, though much more work has been done on the latter than the former.

We therefore agreed that a service level agreement (SLA) driven system design (for composable systems and services) is appropriate, though this raises related issues of the relationship to policy, regulation, and the law. This is further elaborated in the next paragraph. A study of trust boundaries in composed and multi-level systems was also agreed to be important, along with the related policy and legal implications.

The formulation of SLAs needs to be adapted in composed systems, as the probability of a failure grows with the amount of involved parties (as it is intended in composed systems). From a legal viewpoint, it is mandatory to distinguish between external and internal relationships within composed systems. Whereas the external one encompasses the contractual relationship between an end user and the composed system provider, the internal deals with the contractual relationship of the composed system provider and (other) providers of systems s/he himself uses to perform the duties owed to clients. In this situation, the provider of a composed system has different obligations. First, s/he is responsible for the systems s/he himself offers to end users (ensuring that the services being relied on work together as intended). She or he is, in turn, indirectly also responsible for the (contractual) performance of suppliers and their subsystems. These relationships define implicitly a trust boundary already. This boundary can be made explicit with the help of contracts, SLAs, and

technical interfaces that clearly define the owed performance and functionality, and therefore the sphere of responsibility.

However, composed (and in particular virtualized) systems can still cause difficulties with correctly attributing responsibility and liability. Although from a legal viewpoint responsibility lies with the one who acts in a negligent manner in the way that he fails to exercise reasonable care, the complexity of the system itself may make it difficult to pinpoint the cause of a failure (i.e., is it a subsystem or the composition of such systems), and therefore identify the responsible legal party. Monitoring and/or recording systems may be helpful in this situation to assess what actually happened, and permit the party who bears the burden of proof to get an inside view that can support further legal prosecution.

Therefore, the SLAs have to take the higher probability of liability due to the increased number of acting parties into account. Still, it is not obvious that even a well-formulated SLA can procure a non-liability if it is technically not possible to determine what actually happened.

A different question that arose was the following: can we structure (or architect) systems to create boundaries or interfaces that act as trust boundaries, or at least as clear functional or perhaps ownership boundaries. The latter should be relatively easy to achieve, though this immediately reverts to the issue of how to create and agree levels of trust between owners or operators of parts of the infrastructure. A related issue is that of interface abstractions and tussles between entities that are – or that may be – unwilling to exchange information or to agree on trust levels. This is clearly a potential impediment to the successful construction of resilient systems. We simply agreed to add this to our resilient systems research agenda.

Two further, related questions are (i) in what ways are composed systems able to be structured to reduce complexity, where components are not necessarily fully described or understood, and (ii) how to model and understand, and subsequently assure, the resilience of (composed) interconnected and/or interdependent networks. The first of these is studied by complex systems researchers, which was not represented at this Dagstuhl seminar, while the second has recently become the subject of considerable interest amongst graph theorists amongst others, and is recognized by the resilience researchers participating in the seminar as being one of the most important topics for us to study because of the evident interdependence between various real-world (critical infrastructure) networks such as telecommunications, electricity distribution and public transport (for example).

Related to this is that we must ensure that resilience mechanisms do not make systems more fragile, even though we may have made them more complex, and also in introducing them we will very likely have increased the attack surface for the very systems that we are trying to protect.

We moved on to the issue of safety-critical systems, for which there is inevitably different thinking about resilience because of the societal importance of the systems in question, such as aviation, railway transport and roads networks. For these systems, the publication of information about incidents and liability to risks is considered essential and in the public interest. For other networks that can be considered critical (though not safety-critical) such as financial, government, corporate or telecommunication networks, for example, there is much less interest in discussing their resilience, and this is unfortunate when it is increasingly obvious that society really depends on these systems, and they should attract considerably more attention by owners, operators and also researchers.

A key research question, one that is of interest to some participants in the seminar, is understanding and modelling the roles of humans in (composed) systems; this includes how to assess risks, and how to assure resilience of systems in which humans are a constituent

part. Previous research has been conducted in the fields of Human Computer Interface (HCI) and Computer Supported Cooperative Work (CSCW), some time ago, and a current imperative is to study and include this prior research into the resilience research agenda.

We discussed the prospect of autonomic operation in critical systems that need to be made resilient: can removing the human in the loop make safer systems? The range of such systems we discussed included telecommunications operations, aviation, and the driverless car. We observed that in some of these systems, human on the loop (flying on automatic pilot, for example) is already much used. How well would this translate to 'driverless' cars? Clearly there would be significant implications and legal responsibilities and liabilities, which would need to be closely embedded along with discussions about the technical viability about autonomic operation.

And in such composed (especially virtualised) systems there would be difficulties attributing liability (or responsibility), even when activities have been monitored or recorded, following any incidents.

For the goal of assessment it is necessary to find new approaches that can deal with the additional complexity of composed systems, i.e., taking into account new interactions that had not been included in the design and implementation of the individual components. Designing components or systems in the face of uncertainty with respect to their usage context or environment can cause unnecessary complexity in their code to provide the desired level of resilience. This additional complexity caused by uncertain component contexts may weaken the reliability and trust of the overall system, since more possibilities for programming errors exist. Consequently, providing resilience for flexible and composable systems (FCS) is challenging. Approaches like model checking need to be examined whether they can feasibly be applied to this new environment, with an eye towards the on-demand and cost-efficient assessment of the properties of composed systems. Here, also the question of responsibility for these checks and assessments arises, i.e., how much responsibility (and liability) the original components' designers have and how much has to be borne by the composer or end-user. It may be possible to learn more in this respect from other disciplines that have faced/are facing similar issues.

It might be of interest as well to investigate how much a component designer can introduce in terms of mechanisms that are context-sensitive, i.e., change the behaviour of the component in different contexts to maintain SPR. Here, a basic trade-off between high levels of security, privacy and resilience (SPR) on the one hand and a high degree of flexibility/composability on the other needs to be evaluated, as well. One potential solution may be the regular software update of components, which may be necessary or at least desirable for fixing security issues anyway. Such updates may open the possibility to let the component interact in new contexts, but given the longevity of some devices (e.g., sensors built into houses) it is not very likely that vendors provide development and software updates for their products over such long periods.

3.3.2 Question 2

What is the role of law in supporting resilience, privacy [3] and security?

- Does technology regulation play a part?
- How to enforce law – e.g. automated algorithmic law enforcement?

Trying to clarify the potential contribution of law to privacy, security and resilience, one immediately encounters fundamental questions about the function of law in a society, as well as the relation between different approaches to regulation. Entities affected by the law,

particularly companies, expect regulation to be clear, foreseeable, and to provide certainty to enable planning of future decisions. How can these goals be achieved? There is a continuum of norms, from formal laws passed by legislators to standards agreed upon in a technical community. In general, technical standards can react to new developments faster, and given that they are mostly developed by scientists and engineers, they are easier to understand for that same target group. There is sometimes frustration in that community because they seem to be forgotten in the law-making process, so should their voice be considered more when drafting new legislation? This may sometimes improve the quality of the legal texts, but this is difficult to do in a democratic process, which requires elected politicians to be in charge. Neutral advice, of course, would be helpful, but completely neutral experts without their own agenda exist in theory only. The cultural background plays a role as well, as law enforcement alone cannot ensure legal compliance if the law itself is not considered acceptable in society.

One source of frustration when engineers deal with legal texts is the lack of concrete guidelines, which would ideally use precise numbers and thresholds. Unfortunately, this is rarely feasible, as such thresholds would be arbitrary, as seen in examples like a threshold scale from which aerial images in Google Earth are considered to contain personal information. German telecommunications regulation is an example for the inclusion of a concrete regulation model put into law; we doubt whether its interpretation by the Bundesnetzagentur actually follows the spirit of the law, though.

Despite these problems, can the law still play a role in improving security and privacy? We see attempts in Germany (IT Security Act) and the US (particularly in the health sector). Data breach disclosure requirements, for example, can serve as an incentive to improve security, without trying to go into too many details. Previous attempts of regulation in the IT sector have led to unintended results, though, in the broader context, the Oracle vs. Google case is a good example. As an example of infosec regulation that did not work, in the 1980s and 1990s, the US government tried to restrict dissemination of knowledge of public-key encryption and of strong symmetric cryptography, primarily via export control. That led to controversy and "interesting times"; but it is probably not controversial to say that no one's aims were achieved. (Steve Levy's book Crypto is a good summary of this story, but other references exist as well.)

The US NIST's work in establishing cryptographic standards provides both positive and negative examples of effective infosec regulation. One negative example is the process by which NIST and NSA transformed IBM's "Lucifer" algorithm into the Data Encryption Standard (DES); a prevailing belief was that NSA had installed a backdoor. In response, NIST used an open and public competition to select DES's replacement as AES, which consequently had broader acceptance. However, some subsequent actions on the SHA-3 competition raised concerns of backdoors again (as did Snowden's revelations about backdoors in some elliptic curve PRNGs).

A final problem to be considered in information security regulation is the conflict between different jurisdictions, both on the level of federal states and between nations. The concept of discovery, for example, which is used in the U.S., seems rather scary for European lawyers. The above-mentioned export regulation of cryptography is another example.

We also discussed the question how law and technology relate and how they contribute to privacy.

On one hand, law defines a set of rules that are supposed to determine what (among others) technology is allowed to do; on the other hand, we have the impression that legal norms seem to lag behind technical developments. This has led to norms being ignored, and national authorities have often failed to enforce them. In recent years, however, we observe a changing attitude of regulators and courts in Europe.

Common law and civil law jurisdictions have different approaches to cope with technical change. In common law, the focus is on precedents, while civil law uses more abstract norms, trying to cover future cases already in statute (though still requiring a neutral entity, i.e., a court, for arbitration). We assume that the latter approach is better suited to deal with innovation, such as the big data paradigm. European legislation protects personal data, making anonymization a core concept to enable law-compliant data processing|in the big data context, de-anonymization is often possible due to an unforeseeable amount of additional information that can be linked with the original data collection. This leads to the issue when to consider a certain anonymization procedure (such as the addition of noise) as sufficient [4].

In addition, it raises the question whether the current data protection legislation in Europe needs to be adapted to benefit from the advantages of research based on big data. One approach could be to regulate procedure (for data processing) instead of result. In practice, globalization has caused problems for legislation and law enforcement, increasing complexity and enabling circumvention; yet, it does not imply that law is powerless per se. Past experience has shown that the EU has been able to enforce European law even against the interests of global players like Microsoft.

Concerning resilience, we have discussed the impact of (de-)centralization; decentralized systems can increase resilience, but under certain circumstances, the opposite can be true. In networks, users can often change the behaviour of individual nodes, thus causing an impact on the overall system's behaviour that cannot be foreseen or controlled by the respective user or a central entity. The example of the smart grid shows that lack of resilience in IT systems can have real-world consequences. It also illustrates how privacy, security and resilience relate to each other's mechanisms improving privacy, for example, may hinder the detection of attacks, and cryptographic processing may enable DoS attacks due to high processing load.

The group has also covered the topic of trust in IT systems; under which circumstances does one need trust, and can we talk of trust if we are certain of the functionality? Snowden's revelations have, in some cases, shaken the confidence in some beliefs about IT security, thus increasing the relevance of this question once more.

Additional questions/issues which evolved in post-seminar-discussions: What are values law should encompass? What goods should law protect in the field of security and resilience, e.g. ownership and/or personality? Can law and regulatory powers be used to equalize potential market failure and/or differences in power and strength?

3.3.3 Question 3

Traceability of (personal and non-personal) data in service provision?

- Which legal and technical aspects to consider?
- Anonymization and de-anonymization of data?

Users of current and future applications and services (e.g. Gmail, Facebook, Body-Sensors, Smart-homes, etc.) must be assured that the data defining their identity remains protected (and being purpose bound). This requires a new legal definition of data ownership, because the state of the art treats data as tradable goods in a classic sense. Treating this kind of data in this sense is not appropriate since it pertains to a person and defines (directly or indirectly) their identity beyond the transaction period. Hence, it directly links to human rights like the right for privacy and in the case of misuse can ultimately violate the constitutional requirement of keeping the dignity of man sacrosanct. But personal-data still has value that should be exploitable (within a given framework). This value can be tangible (e.g.

expressed through monetary transactions) or intangible (e.g. social standing and personal reputation). In order to facilitate this, a digital "market place" is required that guarantees transaction transparency, awareness and control of personal data when being handled and exploited. Hence personal data cannot be "owned" in the traditional sense but can only be allowed to use within a well-defined legal, political, social and commercial context. In order to achieve this we believe that new standards as regulatory mechanisms are required, which should be legitimated through a democratic process. Research will have to establish how the new concepts of "market place" and "usage rights" can be realized within a technical framework that enables traceability of data and its usage and protects it from misuse (such as unauthorised trading) but still allowing for the realization of new market concepts.

This approach can be extend to systems which would be required to interoperate with each other in dynamic use cases where some may be operated by third parties on behalf of a user which may change over time. A market place approach as above described would already allow management of identity of data. An extension regarding the mapping of resources and services to such identities may allow for management of responsibilities in dynamic unforeseen situations. Future scenarios may include multiple Google-smart-home instances, interoperating with some e-health application.

Additional questions/issues which evolved in post-seminar-discussions: What are the different concepts of identity in law and in technical sciences? What are incentives of the involved parties to accept legal restrictions and how can they be created?

3.3.4 Question 4

How can we improve the perception of assurance [5], privacy, security and resilience for the end-user?

- Can we build an economic model for security and resilience?
- How to specify the role of trust in assuring security?

To ensure security and resilience of (distributed) flexible and composable systems (FCS), in an ideal world, security researchers would be able to test in advance every piece of software or application for potential problems it may encounter, which would allow them to mitigate the problems through adequately modifying the system design. This being overly optimistic, we could still hope for identifying problems as they occur and give adequate warnings to stakeholders, who would then take all and only those actions required to mitigate the risk/damage. Reality often looks different, with information about risks either not available or not distributed timely, or warnings exaggerated and resulting in panic rather than measured response. While not a new problem, FCS is likely to increase the seriousness of these issues.

Regarding FCS, two aspects are particularly challenging: how to measure or assess security, privacy and resilience (SPR), and how to communicate the results of this assessment to the end user of the systems (or other relevant stakeholders). Since assessment is discussed separately, we here focus more on the communication and transparent aspect. Regardless, in both of these aspects one of the main challenges is the fact that, by definition, flexible systems are operated in varying and thus often unforeseeable contexts, e.g., in new service compositions, in new environments, or for new purposes. This characteristic makes it particularly difficult to analyse such systems a priori, and in principle necessitates mechanisms that differ from existing approaches designed for static systems. It may be possible to learn more in this respect from other disciplines that have faced/are facing similar issues. Research in medical drugs could be one such comparator: it is one thing to establish if a drug, taken on its own,

is harmful for a patient. However, it is impossible to foresee what other medication(s) a patient may be taking in addition, outside the confines of a tightly controlled medical trial. There are however mechanisms that with varying success try to address this issue, from voluntary or mandatory reporting mechanisms if an incompatibility has been experienced, to the information leaflet that informs the patient on what other drugs she should avoid. Some of this is underpinned by a risk management strategy. For more serious drugs, pharmacists will ask a set of questions on other medication taken before releasing it to the buyer, with less risky drugs it is left to the patient etc. A possible research question should look at the success or failure of these approaches in cognate fields, and explore to what degree the analogy to FCSs is valid.

The necessity for regulation of such assessment and information mechanisms will depend on a classification of the severity and impact of issues and dangers. Critical, high impact systems with the potential for severe damage (e.g., energy utility systems) should be treated differently than systems that might only have individual, personal effects (although a cumulative effect might be taken into account if these minor damages occur for a large set of end-users). However, social expectations on what counts as "acceptable risk" are changing as rapidly as the technologies themselves. A few years ago, there was no Facebook. Today, even a very short temporary outage leaves people at the very least in emotional distress – some however face real difficulties, as they rely on Facebook to sign into other, more crucial, systems. Here, social practices can cause a loss in resilience that is difficult to foresee or counteract in a FCS environment.

The question of propagating the results to end-users is somewhat dependent on the possibilities for and outcomes of assessment mechanisms. It relates to the concept of trust and how the perception of trust can be enhanced. One basic question is how detailed information about the SPR levels is to be made available to end-users (e.g., a 'five star' system similar to car safety tests, combined abstract and more detailed information like existing energy efficiency classes, or even more detailed reports).

Another is how security breaches or similar critical events are communicated to the parties affected. In the past, this happened largely on an ad-hoc basis, with little or patchy regulation. A conventional virus checker for instance will give some warnings instantaneously, but because this requires automated detection and notification, the untrained user is given relatively little in information on "what to do now". By contrast, a security breach in a credit card company will be communicated, if at all, to customers through the established media with a degree of time delay (or individually, by email or similar channel), but with the advantage of careful advice tailor made to the situation ("it was a minor breach, it is unlikely that your credit card details were released, however you should change your password for this site. If you have further concerns, cancel the card or call our helpline. . . ").

For a world of FCS, we should rethink if these channels of communication are still adequate (if they ever were). With automated bug reporting for instance, to which party should a report be sent if the problem emerged from the ad-hoc interaction between several machines and programs? If we surround ourselves with gadgets and get too many warnings, is there a danger of the "boy who shouted wolf" syndrome, so that we get desensitized and stop taking appropriate actions? Should several machines negotiate with each other which one has to inform a user of an issue, so as not to cause information overload (e.g.: my fridge, thermostat, washing machine and car decide between each other which one to alert me if a problem they all experience has a single cause).

Other connected questions are which level of detail is to be mandatory for selected systems (again linked to the risk classes described above), and to which degree this information needs

to be tailored to the expertise level of the recipient and be made easily accessible. Related to this is investigating the need to educate users about the significance of security information, particularly the 'digital natives', the younger people growing up being used to IT technology and maybe not sensitive to security and privacy issues.

Interesting aspects here might also be the evaluation of the feasibility of wisdom of the crowds approaches (e.g., a futures market for possible attacks, similar to the "futures market in terrorist attacks" that the US government briefly contemplated), or to what degree the existing biological or psychological models for trust are of value in the context of FCS – can we find design solutions that use our evolved mechanisms to ascribe or revoke trust and optimize them?

Apart from legally mandatory information disclosure about critical system characteristics, some of the tasks of informing the end users might also be taken over by market mechanisms. For example, a service or system provider offering more transparency for each customer may be more trustworthy and therefore get more customers. If a sufficient market landscape of service and system operators exists, each with its own approach to transparency in addition to the mandatory one, user preference might lead to the automatic establishment of standards of information. Here relevant research questions should address the drivers and obstacles for an efficient market in FCS – Intellectual property rights (de-jure monopolies) or standard setting for instance could prevent the emergence of an efficient "market in security".

Another market-related incentive for service providers to provide information about its security is if a model for insuring such services can be developed where having higher levels of security and transparency, at least from the viewpoint of the insurance companies, directly results in benefits for the insured provider. In such a scenario, being able to prove that a system is secure would bring direct economic benefits, e.g., being able to show operation without security incidents or usage of more secure systems and receiving a decrease in premiums. On the other hand, insurance companies might also be able to exert pressure by refusing coverage for services where no sufficient transparency is provided. Reliable mechanisms for transparency and auditing are again necessary and useful to this end. An open issue in this context are how feasible such a model is considering the possibility of presently unknown risks such as large scale vulnerabilities and exploits discovered in the future for any service.

Learning again from risk management in other fields could be of benefit. In the regulation of financial services, strict rules exist on what (and when) investors need to be informed, e.g., in the form of a "profit warning". At the same time, individual investors in the UK e.g. will get from their independent financial advisor or investment broker a mandatory "risk profile" that tells them what products are suitable for them given their willingness (and resilience to) certain risks. Could something similar, in machine readable and transportable format be relevant for FCS?

These approaches rely on a final decision on trustworthiness by a human and thus might be of limited use in the important field of application of machine-to-machine communication or automated composition of systems. It is to be seen whether in environments where there are no human intervention or decision other mechanisms are necessary, making trust understandable and utilizable for machines, or if maybe the concept of trust can only be applied for systems where humans are involved. In a world where machines or devices may automatically close contracts, automated trust assessment and delegation may be necessary. One possibility could be developing a machine behaviour code describing acceptable machine behaviours related to an automatically closed contract. This in combination with code breaching detection mechanisms would be an additional component helping to increase SPR.

However, trust and reputation systems tend to become complex and may serve as an attractive attack target instead of attacking other security mechanisms (e.g., attacking a TLS-secured communication may be easier by using illegitimate certificates). On the other hand, removing humans and thus the potential for human error from the loop might also have benefits if reliable automated systems can be designed. In any case, the legal implications of and responsibility attribution in a pure machine-to-machine environment also need to be explored.

It should be noted that an increased importance of assessment and information systems and reliance on the information thus propagated also increases the attractiveness of such meta-systems themselves for attacks. Care needs to be taken that no new avenues of attack are created by mechanisms that themselves are used to assure end users of the resilience, security or privacy of other systems. For instance, a competitor might try to exploit automated fault monitoring and reporting systems by flooding them with faked incidences of faults in a competitor's product (similar to the manipulation of reviews and ratings that we already find on e-commerce and recommendation sites) – there is also a question of how the law should proscribe, if at all, this type of behaviour and impose (criminal law?) sanctions.

Additional questions/issues which evolved in post-seminar-discussions: What role can other actors not directly involved, e.g. insurances, play? Can trust be developed in purely non-human interaction?

3.3.5 Question 5

What constitutes a security problem?

- How to model, define, measure, ... security problems?
- Do we need an ethical framework for handling findings from security research?

It is necessary to inquire into the currently existing and maybe insufficient research infrastructure and culture regarding especially security research. Intellectual property law and data protection law for instance have been accused of hampering necessary security research. While EU data protection law recognizes a "research exemption", there is at least some evidence that this provision is badly understood and insufficient in allying the fears of administrators in university ethics committees. At the very least, the question should be asked if this type of provision that was tailor made for medical research "fits" the practice of security research in FCS. One possibility would be to clarify (or create) "research exceptions" in copyright and data protection law. A potentially more appealing solution would be to restrict these exceptions to a special class of "bona fide security researchers" with additional exemption from legal prosecution for their type of research. Here, lessons could be learned from the very different way the EU and the US regulate journalism as a profession that is also (partially) exempted from data protection rules. At the same time, a code of conduct needs to be established for this research, in particular for rules for the publication of research results that might lead to an increased risk due to the disclosure of security flaws. Another example would be the necessity to report findings independently of their potential to attract attention, i.e., reports of negative results in the sense of not reporting any flaws should be treated the same way as reports about vulnerabilities deemed more 'interesting' for the public. Such a research culture would have as its goal a more thorough and independent research that is thus also perceived as more trustworthy by the general public and by the subjects of investigation. Learning again from the experience with safety research in medicine, it might be worth exploring if there ought to be a "notification scheme" for certain types of research projects and repositories for research findings, to prevent the "publication bias" inherent

in traditional research. In the US, the department for Homeland Security recently made available huge datasets for security research in IT infrastructures through the PREDICT repository. A promising research project would be to evaluate the suitability of this database for research in FCS, and setting up, if needed, a similar system for FCS. The PREDICT approach to data privacy would need to be analysed to ensure its acceptability within an EU setting.

3.3.6 Question 6

How to deal with unforeseen new context of usage?

- What legal and technical dimensions are involved?
- What to expect when shifting responsibility for data?

Even though "How to deal with data usage in a new context" was identified as a stand-alone research topic for a research agenda, we believe that the contribution to topic 3 "Traceability of (personal and non-personal) data in service provision (anonymisation and de-anonymisation of data)" addresses the topic perfectly well.

Additional Material

Original presentations including introductory talks and supporting presentations can be found at theses URLs:

- http://materials.dagstuhl.de/files/15/15151/15151.SWM2.Slides1.ppt
- http://materials.dagstuhl.de/files/15/15151/15151.SWM3.Slides.pptx
- http://materials.dagstuhl.de/files/15/15151/15151.SWM4.Slides.pptx
- http://materials.dagstuhl.de/files/15/15151/15151.SWM5.Slides.pptx
- http://materials.dagstuhl.de/files/15/15151/15151.JerkerDelsing.Slides.pdf
- http://materials.dagstuhl.de/files/15/15151/15151.SaleemBhatti.Slides.pdf

References

1 Rumen Kyusakov, Pablo Punal Pereira, Jens Eliasson, and Jerker Delsing. Exip: a framework for embedded web development. *ACM Transactions on the Web (TWEB)*, 8(4):23, 2014.

2 S. N. Bhatti, D. Phoomikiatissak, and R. J. Atkinson. Fast, Secure Failover for IP. In *MILCOM 2014 – 33rd IEEE Military Communications Conf.*, Oct 2014.

3 Burkhard Schafer. All changed, changed utterly? *Datenschutz und Datensicherheit – DuD*, 35(9):634–638, 2011.

4 Burkhard Schafer, Judith Rauhofer, Zbigniew Kwecka, and William Buchanan. "I am Spartacus": privacy enhancing technologies, collaborative obfuscation and privacy as a public good. *Artificial Intelligence and Law*, 22:113–139, 2014.

5 Aleksandar Hudic, Markus Tauber, Thomas Lorunser, Maria Krotsiani, George Spanoudakis, Andreas Mauthe, and Edgar R. Weippl. A multi-layer and multitenant cloud assurance evaluation methodology. In *Cloud Computing Technology and Science (CloudCom), 2014 IEEE 6th International Conference on*, pages 386–393. IEEE, 2014.

Participants

- Ali Alshawish
Universität Passau, DE
- Silvia Balaban
KIT – Karlsruher Institut für Technologie, DE
- Saleem Bhatti
University of St. Andrews, GB
- Roland Bless
KIT – Karlsruher Institut für Technologie, DE
- Marcus Brunner
Swisscom AG – Bern, CH
- György Dan
KTH Royal Institute of Technology, SE
- Jerker Delsing
Luleå Univ. of Technology, SE
- Thilo Ewald
Microsoft Deutschland GmbH – Unterschleissheim, DE
- Andreas Fischer
Universität Passau, DE
- David Hutchison
Lancaster University, GB
- Youki Kadobayashi
Nara Institute of Science and Technology, JP
- Graham Kirby
University of St. Andrews, GB
- Helmut Leopold
AIT Austrian Institute of Technology – Wien, AT
- Andreas Mauthe
Lancaster University, GB
- Simon Oechsner
NEC Laboratories Europe – Heidelberg, DE
- Frank Pallas
KIT – Karlsruher Institut für Technologie, DE
- Oliver Raabe
KIT – Karlsruher Institut für Technologie, DE
- Alberto Egon Schaeffer-Filho
Federal University of Rio Grande do Sul, BR
- Burkhard Schafer
University of Edinburgh, GB
- Marcus Schöller
Hochschule Reutlingen, DE
- Sean W. Smith
Dartmouth College – Hanover, US
- Christoph Sorge
Universität des Saarlandes – Saarbrücken, DE
- Indra Spiecker gen. Döhmann
Goethe-Univ. Frankfurt, DE
- James P. G. Sterbenz
University of Kansas, US
- Burkhard Stiller
Universität Zürich, CH
- Markus Tauber
AIT Austrian Institute of Technology – Wien, AT
- Gene Tsudik
Univ. of California – Irvine, US
- Pal Varga
Budapest University of Technology & Economics, HU
- Edgar R. Weippl
Secure Business Austria Research, AT