Report from Dagstuhl Seminar 12381

Privacy-Oriented Cryptography

Edited by

Jan Camenisch[1], Mark Manulis[2], Gene Tsudik[3], and Rebecca N. Wright[4]

1 **IBM Research – Zürich, CH,** `jca@zurich.ibm.com`
2 **University of Surrey, GB,** `mark@manulis.eu`
3 **University of California – Irvine, US,** `gts@ics.uci.edu`
4 **Rutgers University, US,** `rebecca.wright@rutgers.edu`

Abstract

This report documents the program of the Dagstuhl Seminar 12381 "Privacy-Oriented Cryptography", which took place at Schloss Dagstuhl in September 16-21, 2012. Being the first Dagstuhl seminar that explicitly aimed to combine cryptography and privacy research communities, it attracted a high number of participants, many of whom were new to Dagstuhl. In total, the seminar was attended by 39 international researchers, working in different areas of cryptography and privacy, from academia, industry, and governmental organizations. The seminar included many interactive talks on novel, so-far unpublished results, aiming at the design, analysis, and practical deployment of cryptographic mechanisms for protecting privacy of users and data. The seminar featured two panel discussions to address various approaches towards provable privacy and different challenges but also success stories for practical deployment of existing cryptographic privacy-oriented techniques.

Seminar 16.–21. September, 2012 – www.dagstuhl.de/12381
1998 ACM Subject Classification D.4.6 Security and Protection
Keywords and phrases Privacy, Cryptography, Anonymity, Confidentiality
Digital Object Identifier 10.4230/DagRep.2.9.165

1 Executive Summary

Jan Camenisch
Mark Manulis
Gene Tsudik
Rebecca N. Wright

The constantly increasing volume of electronic interactions and sensitive information disseminated online raises privacy concerns and motivates the need for efficient privacy-oriented techniques. The aim of our "Privacy-Oriented Cryptography" seminar was to bring together (mainly, but not only) researchers working in different domains of cryptography and privacy. Although non-cryptographic measures can, at times, aid privacy (e.g., statistical or ad hoc obfuscation techniques) — cryptography, via its mathematical mechanisms and formal concepts, helps obtain novel and efficient privacy-enhancing solutions, achieving concrete and measurable privacy guarantees.

Since privacy is a very broad area, being explored not only by security and cryptography experts, this seminar focused on two domains: *user privacy* and *data privacy*, for which the

Privacy-Oriented Cryptography, *Dagstuhl Reports*, Vol. 2, Issue 9, pp. 165–183
Editors: Jan Camenisch, Mark Manulis, Gene Tsudik, and Rebecca Wright
Dagstuhl Reports
Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

benefit from using cryptographic techniques is especially significant. Seminar participants presented and discussed many novel privacy-oriented cryptographic algorithms and protocols that admit various fields of deployment for protecting privacy in a broad range of applications, involving possibly huge amounts of data (e.g., cloud computing) and many different users (e.g. online communities). The seminar further addressed the emerging research direction of *provable privacy*, by discussing various mechanisms and techniques for defining concrete privacy goals and enabling their formal analysis.

The seminar brought together 39 of the leading scientists in the areas of (applied) cryptography and privacy. The participants came from all over the world, including the US (13 participants), Germany (8), Switzerland (6), Great Britain (5), Australia (1), Belgium (1), Canada (1), France (1), Italy (1), and Sweden (1).

The program contained 26 interactive presentations, each about 35–40 minutes and two panel discussions, with a free afternoon on Wednesday to offer time for social activities or for conducting collaborative research in smaller groups. The seminar ended on Friday after lunch to enable time for traveling. We asked participants prior to the seminar to suggest talks based on their most recent results. Most presentations followed this suggestion and introduced new, sometimes even not yet submitted or still work-in-progress results. The first panel — “Privacy Models: UC or Not UC?” — discussed the advantages and disadvantages of existing cryptographic methods for formal specification and analysis of security and privacy guarantees. The second panel — “Privacy-Oriented Cryptography: Why is it not adopted more in practice?” — discussed challenges that arise in the practical deployment of existing privacy-oriented cryptographic solutions but also considered some success stories like Tor, a popular anonymous communications service, which is widely used in different parts of the world.

The nature of the seminar allowed experts and practitioners to air ideas and discuss preliminary concepts and work-in-progress results. This might have led to the exposure and subsequent exploration of new research directions that may offer both practical significance and intellectual challenge.

The organizers would like to thank all participants for accepting our invitations and attending the seminar, and for sharing their ideas and contributing to the interesting seminar program. We hope that discussions were fruitful and the opportunity to work face-to-face during the seminar helped to create impulses for exciting new research projects, paving the way for further progress and new discoveries in Privacy-Oriented Cryptography.

Finally, the organizers, also on behalf of the participants, would like to thank the staff and the management of Schloss Dagstuhl for their support throughout the 1,5 years of preparations of this very pleasant and successful event.

2 Table of Contents

Executive Summary

Overview of Talks

3 Overview of Talks

3.1 Hacking Smart Machines with Smarter Ones: How to Extract Meaningful Data from Machine Learning Classifiers

Giuseppe Ateniese (Johns Hopkins University – Baltimore, US)

Joint work of Ateniese, Giuseppe; Felici, Giovanni; Mancini, Luigi V.; Spognardi, Angelo; Villani, Antonio; Vitali, Domenico

Machine Learning (ML) algorithms are used to train computers to perform a variety of complex tasks and improve with experience. Computers learn how to recognize patterns, make unintended decisions, or react to a dynamic environment. Certain trained machines may be more effective than others because they are based on more suitable ML algorithms or because they were trained through superior training sets. The distinctive ability of a particular ML model gradually develops during the training phase, when it is fed with samples of data to form its knowledge. Although ML algorithms are known and publicly released, training sets may not be reasonably ascertainable and, indeed, may be guarded as trade secrets.

While much research has been performed about the privacy of the elements of training sets, in this paper we focus our attention on ML classifiers and on the statistical information that can be unconsciously or maliciously revealed from them. We show that it is possible to infer unexpected but useful information from ML classifiers.

In particular, we build a novel meta-classifier and train it to hack other classifiers, obtaining meaningful information about their training sets. This kind of information leakage can be exploited, for example, by a vendor to build more effective classifiers or to simply acquire trade secrets from a competitor's apparatus, potentially violating its intellectual property rights.

3.2 Attribute-Based Encryption from Lattices

Xavier Boyen (Prime Cryptography – Palo Alto, US)

We initiate the study of a new promising framework for the design of expressive cryptosystems from lattice assumptions. Specifically, we construct the first ever “complex” ABE scheme for a post-quantum world, defeating critical obstacles previously standing in the way of this actively researched result.

3.3 Enabling Complex Queries and RBAC Policies for Multiuser Encrypted Storage

Bruno Crispo (University of Trento – Povo, IT)

Cloud computing has the advantage that it offers companies (virtually) unlimited data storage at attractive costs. However, it also introduces new challenges for protecting the confidentiality of the data. Most current security schemes support an all-or-nothing access model to the data that is too coarse-grained. Moreover, existing schemes do not allow complex encrypted queries over encrypted data in a multi-user setting. Instead, they are limited to keyword searches or conjunctions of keywords.

We extend the work done on multi-user encrypted search schemes by (i) supporting SQL-like encrypted queries on encrypted databases, and (ii) introducing a fine-grained access control over the data stored in the outsourced database based on the RBAC model.

Finally, we implemented our scheme and compare its performance with recent similar solutions.

3.4 Efficient and Secure Testing of Fully-Sequenced Human Genomes

Emiliano De Cristofaro (PARC – Palo Alto, US)

Joint work of Baldi, Pierre; Baronio, Roberta; De Cristofaro, Emiliano; Gasti, Paolo; Tsudik, Gene
Main reference P. Baldi, R. Baronio, E. De Cristofaro, P. Gasti, G Tsudik, “Countering GATTACA: efficient and secure testing of fully-sequenced human genomes,” in Proc. of the 18th ACM Conf. on Computer and Communications Security (CCS'11), pp. 691–702, ACM, 2011.
URL http://dx.doi.org/10.1145/2046707.2046785

Recent advances in DNA sequencing technologies have put ubiquitous availability of fully sequenced human genomes within reach. It is no longer hard to imagine the day when everyone will have the means to obtain and store one's own DNA sequence. Widespread and affordable availability of fully sequenced genomes immediately opens up important opportunities in a number of health-related fields. In particular, common genomic applications and tests performed in vitro today will soon be conducted computationally, using digitized genomes. New applications will be developed as genome-enabled medicine becomes increasingly preventive and personalized. However, this progress also prompts significant privacy challenges associated with potential loss, theft, or misuse of genomic data.

In this talk, we begin to address genomic privacy by focusing on some important applications: Paternity Tests, Ancestry Testing, Personalized Medicine, and Genetic Compatibility Tests. After carefully analyzing these applications and their privacy requirements, we propose a set of efficient techniques based on private set operations. This allows us to implement in silico some operations that are currently performed via in vitro methods, in a secure fashion. Experimental results demonstrate that proposed techniques are both feasible and practical today.

Finally, we explore a few alternatives to securely store human genomes and allow authorized parties to run tests such that only the required minimum amount of information is disclosed.

3.5 On Anonymity Attacks in the Real-World

Roger Dingledine (The TOR Project, US)

Tor's approach to threat models is to try to understand the capabilities of realistic attackers we expect to encounter, rather than picking adversaries our protocols can withstand. This strategy has led us to deploy systems that are not amenable to security proofs. Or to say it even more strongly, we deploy provably insecure systems relative to real-world adversaries, because they're still the safest ones we can deploy.

In this talk I'll explain some realistic attacks against Tor's anonymity and blocking-resistance properties, and discuss some reasons why it's hard to produce accurate and useful models for these attacks (and thus hard to prove things about them).

3.6 Overcoming Weak Expectations

Yevgeniy Dodis (New York University, US)

Recently, there has been renewed interest in basing cryptographic primitives on weak secrets, where the only information about the secret is some non-trivial amount of (min-)entropy. From a formal point of view, such results require to upper bound the expectation of some function $f(X)$, where X is a weak source in question. We show an elementary inequality which essentially upper bounds such 'weak expectation' by two terms, the first of which is independent of f, while the second only depends on the 'variance' of f under uniform distribution. Quite remarkably, as relatively simple corollaries of this elementary inequality, we obtain some 'unexpected' results, in several cases noticeably simplifying/improving prior techniques for the same problem. Examples include non-malleable extractors, leakage-resilient symmetric encryption, seed-dependent condensers and improved entropy loss for the leftover hash lemma.

3.7 Cryptographic Protocols for Privacy Preserving Access Control in Databases

Maria Dubovitskaya (IBM Research – Zürich, CH)

Joint work of Camenisch, Jan; Dubovitskaya, Maria; Neven, Gregory

We present cryptographic protocols that allow querying a database in a privacy preserving way under access control restrictions. We first design a protocol for the oblivious transfer of data with access control mechanisms. This work can be extended for a practical application of buying records from a database privately with unlinkable priced oblivious transfer protocol. Other extensions use attribute-based encryption and anonymous credentials for obliviously querying a database with hidden access control policies of different complexity.

3.8 Privacy Concepts in the New German Electronic Identity Cards

Marc Fischlin (TU Darmstadt, DE)

We review the privacy aspects of the protocols in the new German electronic identity cards, which have been issued since November 2011. Some of the protocols have also been adopted by the International Civil Aviation Organization (ICAO) and can be expected to be deployed for international machine readable travel documents.

3.9 ZQL: A Cryptographic Compiler for Processing Private Data

Cedric Fournet (Microsoft Research UK – Cambridge, GB)

Joint work of Danezis, George; Fournet, Cedric; Kohlweiss, Markulf; Luo, Zhengqin

Our goal is to enable programming on sensitive data without disclosing it. To this end, we compile SQL queries to be executed on client-side datasets, and automatically produce protocols that guarantee both integrity for the query results and confidentiality for the rest of the data. Our protocols rely on zero- knowledge cryptographic evidence, so that query evaluations can be checked without leaking information. We have built prototype compilers that produce C and F#; the F# code can be verified on the fly using our latest type systems for security. We illustrate our approach on queries for paying utility bills and pay-as-you-go insurance policies based on the detailed readings provided by smart meters and GPS units.

3.10 All-but-k Commitments

Ian Goldberg (University of Waterloo, CA)

Joint work of Henry, Ryan; Goldberg, Ian; Kate, Aniket; Zaverucha, Gregory; Olumofin, Femi; Huang, Yizhou

In this talk, we present a new kind of commitment scheme called "all-but- k commitments". With these commitments, Alice can commit to $n - k$ items, and then later open the commitment to n items. Bob is assured that Alice knew $n - k$ of those n items at the time of the commitment, but does not know which. We demonstrate an efficient implementation of all-but-k commitments using Kate et al.'s polynomial commitments from Asiacrypt 2010.

We illustrate the need for all-but-k commitments by demonstrating an attack on the soundness of Peng and Bao's "batch zero-knowledge proof and verification" protocol for proving knowledge and equality of one-out-of- n pairs of discrete logarithms. We repair the protocol using our new commitment construction, and in fact can easily generalize the repaired protocol to a k-out-of-n setting. For $k = 1$, this yields an "OR" proof; for $k = n$, this yields an "AND" proof. For intermediate values of k, this batch protocol is entirely novel.

3.11 Tightly Secure Signatures and Public-Key Encryption

Dennis Hofheinz (KIT – Karlsruhe Institute of Technology, DE)

Joint work of Hofheinz, Dennis; Jager, Tibor
Main reference D. Hofheinz, T. Jager, "Tightly Secure Signatures and Public-Key Encryption," in Advances in Cryptology (CRYPTO'12), LNCS, vol. 7417, pp. 590–607, Springer, 2011.
URL http://dx.doi.org/10.1007/978-3-642-32009-5_35

We construct the first public-key encryption scheme whose chosen- ciphertext (i.e., IND-CCA) security can be proved under a standard assumption and does not degrade in either the number of users or the number of ciphertexts. In particular, our scheme can be safely deployed in unknown settings in which no a-priori bound on the number of encryptions and/or users is known.

As a central technical building block, we devise the first structure-preserving signature scheme with a tight security reduction. (This signature scheme may be of independent interest.) Combining this scheme with Groth-Sahai proofs yields a tightly simulation-sound non-interactive zero-knowledge proof system for group equations. If we use this proof system in the Naor- Yung double encryption scheme, we obtain a tightly IND-CCA secure public-key encryption scheme from the Decision Linear assumption.

We point out that our techniques are not specific to public-key encryption security. Rather, we view our signature scheme and proof system as general building blocks that can help to achieve a tight security reduction.

3.12 On Genomic Privacy

Jean Pierre Hubaux (EPFL – Lausanne, CH)

Joint work of Ayday, E.; Fellay, J.; Hubaux, J.-P.; Humbert, M.; McAllen, P.; Raisaro, J. L.; Rougemont, J.

Over the last 20 years, DNA sequencing has progressed faster than Moore's law. This amazing progress is paving the way to so-called "personalized medicine", meaning that treatments can be fine-tuned to the genetic profile of each patient. In some hospitals, systematic sequencing of (consenting) patients is expected to happen as early as next year.

Yet, this evolution has dramatic implications in terms of privacy. If they happened to fall in the wrong hands, genomic data could be used to perpetrate various kinds of social discrimination (e.g. for decisions related to health/life insurance, recruitment, mortgage,...). Such data can also be used for paternity tests.

In this talk, we will provide an introduction to genomics for computer scientists and briefly describe the (few) ongoing research efforts in this field. We will also explain our own research project, focused on disease- susceptibility tests. Homomorphic encryption makes it possible to preserve diagnostic accuracy while protecting data privacy, at a reasonable cost.

Finally, we will show that the research challenges are formidable and that much more work needs to be done to meet the privacy challenges raised by genomics and more generally by the expected fundamental transformation of medicine [1].

References

1 E. Topol. *The Creative Destruction of Medicine.* Basic Books, 2012

3.13 LIRA: Lightweight Incentivized Routing for Anonymity

Aaron Johnson (Naval Research – Washington, US)

Joint work of Jansen, Rob; Johnson, Aaron; Syverson, Paul

Tor, the most popular deployed distributed onion routing network, suffers from performance and scalability problems stemming from a lack of incentives for volunteers to contribute. Insufficient capacity limits scalability and harms the anonymity of its users.

We introduce LIRA, a lightweight scheme that creates performance incentives for users to contribute bandwidth resources to the Tor network. LIRA uses a novel cryptographic lottery: winners may be guessed with tunable probability by any user or bought in exchange for resource contributions. The traffic of those winning the lottery is prioritized through Tor. The uncertainty of whether a buyer or a guesser is getting priority improves the anonymity of those purchasing winners, while the performance incentives encourage contribution. LIRA is more lightweight than prior reward schemes that pay for service and provides better anonymity than schemes that simply give priority to traffic originating from fast relays.

We analyze LIRA's efficiency, anonymity, and incentives, present a prototype implementation, and describe experiments that show it indeed improves performance for those servicing the network.

3.14 Towards Automizing Secure Two-Party Computation

Stefan Katzenbeisser (TU Darmstadt, DE)

The practical application of Secure Two-Party Computation is hindered by the difficulty to implement secure computation protocols. While recent work has proposed very simple programming languages which can be used to specify secure computations, it is still difficult for practitioners to use them, and cumbersome to translate existing source code into this format. Similarly, the manual construction of two-party computation protocols, in particular ones based on the approach of garbled circuits, is labor intensive and error-prone.

The talk describes the design of a new tool called CBMC-GC, which achieves Secure Two-Party Computation for ANSI C programs. Our work is based on a combination of model checking techniques and two-party computation based on garbled circuits. Our key insight is a nonstandard use of the bit- precise model checker CBMC which enables us to translate ANSI C programs into equivalent Boolean circuits. To this end, we modify the standard CBMC translation from programs into Boolean formulas whose variables correspond to the memory bits manipulated by the program. We demonstrate that CBMC-GC can compile reasonably-sized programs and achieves practical performance.

3.15 Ideal APIs and Modular Verification for TLS 1.2

Markulf Kohlweiss (Microsoft Research UK – Cambridge, GB)

TLS is possibly the most used secure communications protocol, and also the most studied, with a 18-year history of flaws and fixes, ranging from its protocol logic to its cryptographic design, and from the Internet standards to its numerous implementations. We develop a verified reference implementation of TLS 1.2. Our code fully supports its wire formats, ciphersuites, sessions and connections, re-handshakes and resumptions, alerts and errors, and data fragmentation, as prescribed in the RFCs; it interoperates with mainstream web browsers and servers. At the same time, our code is carefully structured to enable its modular, automated verification, from its main API down to computational assumptions on its sub-protocols and underlying cryptographic algorithms.

Our implementation is written in F# and specified in F7. We present its main interfaces—which may be read as specifications of large ideal functionalities—for its main components, such as authenticated encryption for the record layer and key establishment for the handshake. We describe its verification using the F7 refinement typechecker. To this end, we equip each cryptographic and construction of TLS with a new, ideal typed interface that captures its security properties, and we show how to replace their implementations with ideal functionalities while preserving indistinguishability. We thus obtain precise security results for a large part of TLS, for every byte of every connection, relative to the strength of the handshake and the record layer cryptography as selected by the negotiated ciphersuite. We also report new attacks.

3.16 On The Security of One-Witness Blind Signature Schemes, and on Some Alternatives

Anna Lysyanskaya (Brown University – Providence, US)

Joint work of Baldimtsi, Foteini; Lysyanskaya, Anna

Blind signatures have proved an essential building block for applications that protect privacy while ensuring unforgeability, e.g., electronic cash and electronic voting. One of the oldest, and most efficient blind signature schemes is the one due to Schnorr that is based on his famous identification scheme. Although it was proposed over twenty years ago, its unforgeability remains an open problem, even in the random-oracle model. In this talk, I will first show that current techniques for proving security in the random oracle model do not work for the Schnorr blind signature. Our results generalize to other important blind signatures, such as the one due to Brands. Brands' blind signature is at the heart of Microsoft's UProve system, which makes this work relevant to cryptographic practice as well.

This negative result naturally leads to the next question: Can we achieve the attractive features of UProve (lightweight, albeit linkable, anonymous credentials with attributes from just a couple of exponentiations in elliptic curve groups) in a provably secure fashion? Blind signatures alone do not give us the desired functionality; instead, we define blind signatures *with attributes* and give a construction for those whose efficiency is comparable to that of

UProve, and whose security relies on the decisional Diffie-Hellman assumption.

This talk is based on two joint papers with Foteini Baldimtsi:

– http://eprint.iacr.org/2012/197 and

– http://eprint.iacr.org/2012/298.

3.17 Relationships among Privacy Notions for Signatures

Mark Manulis (University of Surrey, GB)

Joint work of Fleischhacker, Nils; Günther, Felix; Kiefer, Franziskus; Manulis, Mark; Poettering, Bertram
Main reference N. Fleischhacker, F. Günther, F. Kiefer, M. Manulis, B. Poettering, "Pseudorandom Signatures," Cryptology ePrint Archive: Report 2011/673.
URL http://eprint.iacr.org/2011/673

Research on privacy for (ordinary) digital signatures, initiated by Yang et al. (PKC 2006) and continued by Fischlin (PKC 2007) demonstrated that for high-entropy hidden messages digital signatures can provide signer's anonymity. Later, Dent et al. (PKC 2010) showed that in this setting digital signatures can also provide confidentiality for the signed messages.

Building on these results I'll show that in fact digital signatures admit much stronger privacy guarantees than previously thought. It is namely possible to hide the entire information about the signature scheme, including its parameters, specification of algorithms, etc.

I'll talk about the new notion and constructions of *pseudorandom signatures*, which essentially guarantee that no adversary can distinguish between a string that was output by the signing algorithm and some randomly chosen bit string (of appropriate length). I'll demonstrate how this notion relates to existing notions of anonymity and confidentiality and propose efficient techniques that can be used to construct pseudorandom signatures.

This talk is based on a joint paper with Nils Fleischhacker, Felix Günther, Franziskus Kiefer, and Bertram Poettering, http://eprint.iacr.org/2011/673.

3.18 Practical Yet Universally Composable Two-Server Password-Authenticated Secret Sharing

Gregory Neven (IBM Research – Zürich, CH)

Joint work of Camenisch, Jan; Lysyanskaya, Anna; Neven, Gregory

Password-authenticated secret sharing (PASS) schemes, first introduced by Bagherzandi et al. at CCS 2011, allow users to distribute data among several servers so that the data can be recovered using a single human- memorizable password, but no single server (or even no collusion of servers up to a certain size) can mount an off-line dictionary attack on the password or learn anything about the data.

We propose a new, universally composable (UC) security definition for the two-server case (2PASS) in the public-key setting that addresses a number of relevant limitations of the previous, non-UC definition. For example, our definition makes no prior assumptions on the distribution of passwords, preserves security when honest users mistype their passwords, and guarantees secure composition with other protocols in spite of the unavoidable non-negligible success rate of online dictionary attacks. We further present a concrete 2PASS protocol and prove that it meets our definition. Given the strong security guarantees, our protocol is surprisingly efficient: in its most efficient instantiation under the DDH assumption in the random-oracle model, it requires fewer than twenty elliptic-curve exponentiations on the user's device. We achieve our results by careful protocol design and by exclusively focusing on the two-server public-key setting.

3.19 Identifying Common Friends in a Privacy-Preserving Way

Bertram Poettering (RHUL – London, GB)

Joint work of De Cristofaro, Emiliano; Manulis, Mark; Poettering, Bertram
Main reference E. De Cristofaro, M. Manulis, B. Poettering, "Private Discovery of Common Social Contacts," Cryptology ePrint Archive: Report 2011/026
URL http://eprint.iacr.org/2011/026

The past decade witnessed a plethora of novel platforms and techniques for online social interaction, including different forms of online social networks, ubiquitous computing on smartphones, and so on. Clearly these developments also pose novel privacy threats on participants and their data.

This talk focuses on a specific problem that arises when users of a social network want to learn the set of their common friends in a privacy- preserving way, i.e. without disclosing non-matching contacts to each other, and without relying on a trusted third party (which might not be reachable in a ubiquitous environment). We offer a full cryptographic treatment of the problem, including security models and provably-secure solutions.

3.20 Data Privacy at Scale

Bartosz Przydatek (Google Switzerland – Zürich, CH)

As users entrust more and more private data to the cloud, it is critical to provide adequate protections for the data, yet without sacrificing high performance, availability and functionality of services, and without impeding innovation.

I will talk about challenges in protecting data privacy at scale, and will discuss a suite of technologies we developed to help addressing these challenges. In particular, I will touch upon issues caused by the lack of key management infrastructure accessible for the general public, and explain why even an efficient fully homomorphic encryption would probably not solve all the problems. I will describe then a simple scalable architecture in which data encryption with appropriate authorization checks and logging provides a meaningful protection for data privacy.

3.21 New Privacy Issues in UMTS

Jean-Pierre Seifert (TU Berlin, DE)

Joint work of Arapinis, Myrto; Mancini, Loretta; Ritter, Eike; Ryan, Mark; Golde, Nico; Redon, Kevin; Borgaonkar, Ravishankar

Mobile telephony equipment is daily carried by billions of subscribers everywhere they go. Avoiding linkability of subscribers by third parties, and protecting the privacy of those subscribers is one of the goals of mobile telecommunication protocols.

We use formal methods to model and analyse the security properties of 3G protocols. We expose two novel threats to the user privacy in 3G telephony systems, which make it possible to trace and identify mobile telephony subscribers, and we demonstrate the feasibility of a low cost implementation of these attacks. We propose fixes to these privacy issues, which also take into account and solve other privacy attacks known from the literature. We successfully prove that our privacy-friendly fixes satisfy the desired unlinkability and anonymity properties using the automatic verification tool `ProVerif`.

3.22 SSL: Myth or Reality?

Vitaly Shmatikov (University of Texas at Austin, US)

Originally deployed in Web browsers, SSL (Secure Sockets Layer) has become the de facto standard for secure Internet communications and is now used widely even in non-browser software. SSL is intended to provide end-to-end security even against an active, man-in-the-middle attacker.

It turns out that SSL is completely insecure against a man-in-the-middle attack in many critical applications and libraries. Vulnerable software includes Amazon's EC2 Java library and all cloud clients based on it, Amazon's and PayPal's merchant SDKs responsible for transmitting payment details from e-commerce sites to payment gateways, integrated shopping carts such as osCommerce, ZenCart, Ubercart, and PrestaShop, AdMob code used by mobile websites, Chase mobile banking and many other Android apps and libraries, as well as Java Web-services middleware and all applications based on it. Interestingly, all these programs use correct SSL implementations... badly.

I will discuss the root causes of these vulnerabilities and present some recommendations for the developers of SSL libraries and applications that use SSL.

This talk is based on the forthcoming ACM CCS 2012 paper.

3.23 Practical Oblivious Access at 1Mbps+

Radu Sion (Stony Brook University, US)

We review several of our recent results including the first/fastest practical ORAM linux file system as well as other fun cloud-related crypto. This is also a preview of our CCS 2012 ORAM papers.

3.24 Anonymity and One-Way Authentication in Key Exchange Protocols

Douglas Stebila (Queensland University of Technology, AU)

Main reference I. Goldberg, D. Stebila, B. Ustaoglu, "Anonymity and one-way authentication in key exchange protocols," in *Designs, Codes and Cryptography.*
URL http://dx.doi.org/10.1007/s10623-011-9604-z

Key establishment is a crucial cryptographic primitive for building secure communication channels between two parties in a network. It has been studied extensively in theory and widely deployed in practice. In the research literature a typical protocol in the public-key setting aims for key secrecy and mutual authentication. However, there are many important practical scenarios where mutual authentication is undesirable, such as in anonymity networks like Tor, or is difficult to achieve due to insufficient public-key infrastructure at the user level, as is the case on the Internet today.

In this work we are concerned with the scenario where two parties establish a private shared session key, but only one party authenticates to the other; in fact, the unauthenticated party may wish to have strong anonymity guarantees. We present a desirable set of security, authentication, and anonymity goals for this setting and develop a model which captures these properties. Our approach allows for clients to choose among different levels of authentication. We also describe an attack on a previous protocol of Overlier and Syverson, and present a new, efficient key exchange protocol that provides one-way authentication and anonymity.

3.25 Privacy-Preserving Interval Operations

Susanne Wetzel (Stevens Institute of Technology, US)

Joint work of Mayer, Daniel; Meyer, Ulrike; Wetzel, Susanne

In this talk we present some work-in-progress on designing privacy- preserving protocols for operations on intervals of integers. In particular, we will present new 2-party protocols to test for the overlap of two intervals, to determine the size of the overlap of two intervals, and to select a random sub-interval in the overlap. We will show that the new protocols are privacy- preserving in the context of a semi-honest adversary.

3.26 Recent Attacks On Mix-Nets

Douglas Wikstroem (KTH Stockholm, SE)

We revisit mix-nets with randomized partial checking (RPC) as proposed by Jakobsson, Juels, and Rivest (2002). RPC is a technique to verify the correctness of an execution both for Chaumian and homomorphic mix-nets. We identify serious issues in the original description of mix-nets with RPC and show how to exploit these to break both correctness and privacy. Our attacks are practical and applicable to real world mix-net implementations, e.g., the Civitas and the Scantegrity voting systems.

We also consider the heuristically secure mix-net proposed by Puiggalí and Guasch (EVOTE 2010) used in the recent large scale electronic elections in Norway. We present practical attacks on both correctness and privacy for some sets of parameters of the scheme. Currently, we are unable to leverage this into an attack on the electronic election scheme as a whole due to additional components.

4 Working Groups

Many participants took the opportunity offered by Dagstuhl's cosy atmosphere to work face-to-face in smaller groups. This work included continuation of existing research collaborations, engagement in smaller group discussions, and initiation of new research agendas.

5 Open Problems

One of the main open problems, as considered by many seminar participants, is that existing privacy models are often too narrow and do not address various privacy threats that exist in the real world. As a result, development of more sophisticated privacy models and appropriate design of provably private yet practically relevant privacy-oriented security protocols and mechanisms was identified as an important research direction for the coming years.

6 Panel Discussions

The seminar included two panel discussions on the models and definitions of privacy and on the use of privacy-oriented cryptography in practice. These panels discussion are summarized in the following.

6.1 Panel "Privacy Models: UC or Not UC?"

Moderator: Jan Camenisch
Panelists: Giuseppe Ateniese, Yevgeniy Dodis, Ian Goldberg, Dennis Hofheinz

There are a number of different approaches for proving the security of a cryptographic protocol or primitive. The probably most popular ones are game based definitions and the universal composability (UC) model. In the former model, security is defined by a number of games, each capturing one (security) property that the protocol/primitive should satisfy.

This approach typically leads to shorter and easier to read proofs but has the drawback the protocol/primitive might satisfy each property separately but not all of the at the same time. In contrast, in the UC model, security is defined by specifying an ideal process describing the expected behavior of the protocols and then it is proved that the participants of the protocol cannot distinguish whether they interact with the ideal process or the protocol.

The panelist and the audience vividly discusses the topic, not only during the panel but also throughout the whole workshop. There seems to have been a common agreement that specifying a protocol by an ideal process is the right thing to do. It captures precisely what a protocol achieves on the one hand and in principle provides users of cryptography such as system designers with an accessible description of what a protocol achieves and how is can be used. There also seemed to have been agreement that the current UC-style models do not quite achieve this yet for a number of reasons. First, the specifications are often very complex and therefore hard to understand so that still it's not easy for systems designer to employ cryptography. Second, the proofs in the UC-style models are often much much longer that in game based definitions. Thus, the proofs are hard to write and hard to verify (probably very few people apart from the authors read them). The panelists and audience felt that it would be nice if:

- Guidelines for defining an ideal process existed, similar to, e.g., JAVA programming language textbooks.
- That probably a second abstraction layer (or at least explanation) is needed that makes a ideal process specification more accessible, together with a (provable) reduction to the ideal process specification.
- Better tools to structure proofs existed so that they are easier to write and read and therefore more confidence in the correctness of the proofs is achieved.

6.2 Panel "Privacy-Oriented Cryptography: Why is it not adopted more in practice?"

Moderator: Rebecca N. Wright
Panelists: Emiliano de Cristofaro, George Danezis, Anna Lysyankaya, Kai Rannenberg, and Radu Sion
Invited Commenters: Roger Dingledine and Ian Goldberg

Privacy-oriented cryptography can provide enhanced privacy protection for individuals, businesses, and governments in a large variety of situations, including web personalization, smart metering, health care, and surveillance. Beyond just protecting the privacy of information in transit or in storage, privacy-oriented cryptography enables "computing without revealing", in which parties can carry out certain computations or interactions, while only learning specific results. This panel addressed the extent to which privacy-oriented cryptography is and is not adopted for practical use, and barriers that must be overcome in order for it to have increased adoption.

The panelists and other participants noted that in fact there has already been some adoption of privacy-oriented cryptography in practice, citing, for example: Open SSL, which has massive deployment and use; German identity cards, which contain privacy-friendly cryptography; the ISO definition of identity, which is based on attributes only; and the Tor project, which provides anonymous Internet communication. Additionally, some felt that it is simply too early in the development of such technologies and supporting technologies on which they rely (e.g., fast communications and large storage capabilities) to expect widespread adoption. As a comparison, nuclear energy has been around for decades, and it still faces technology problems and resistance. Evidence toward continued progress are that companies are beginning to pay attention to the value of adopting privacy technologies, (e.g. MSIE and DoNotTrack, encryption in the cloud and in telcos), and governments have programs (NIST & Trust identities, DARPA & computation over encrypted data) developing and standardizing these technologies.

The panelists and other participants also discussed factors that potentially limit the practical adoption of privacy-oriented cryptography and how they might be overcome. These include the following:

- Privacy-oriented cryptography solutions are difficult to understand, and cryptographers don't spend enough time explaining their ideas. Some noted that the job of cryptographers is to push the state-of-the-art of designs, and let professional engineers incorporate these into systems in ways that cryptographers can't necessarily predict. Others noted that there is a credibility gap between cryptographers and engineers. To bridge this gap, engineers need to understand the assumptions and tools of cryptography, and cryptographers need to take engineering concerns seriously. This takes willingness, time, and effort by both cryptographers and engineers, and may take decades.
- Researchers do not typically carry the results far enough into the technology pipeline to achieve real-world deployment. This is partly caused by funding and training incentives for researchers to focus on publishing papers and moving on to new results rather than carrying individual ideas through to deployable and deployed systems. Further, it is sometimes easier to publish papers that break designs than that build them. As a result, privacy technology designs often do not consider how end users will understand and interact with them, further limiting their potential for adoption.
- There is a real gap between problems that theoretically seem like privacy-preserving cryptography can solve and the actual realities that are revealed when adoption is attempted. This requires feedback with the community that must respond to new understandings of the challenges to be solved. Even when developed, privacy technology is often difficult to understand and use, even for experts. This means even when the technology is used, it is often used incorrectly.
- People always say they value privacy, but quickly lose interest when they must pay for it, learn how to use new technologies, or otherwise change their behavior.
- Companies are not incentivized to protect consumer privacy, only business privacy. There is a perception (often a misperception, we think) that consumer data is a gold mine and privacy technology and legislation can only hurt profits. Businesses only protect privacy to the extent that their (largely under-informed) customers demand it.
- Privacy legislation may be the best way to drive large-scale deployment of privacy technologies. However, convincing governments to adopt technology or legislation is a political process. This is not something that researchers are typically well-equipped to handle, both by training and by funding.
- Free technology is more likely to be adopted than technologies that individuals or companies must pay for, but these must be paid for somehow. Example successes paid for by government funding and volunteer time of an implementation/adoption community, motivated by keeping users safe: Open SSL, Off-the-Record messaging, Tor.

Participants

- Giuseppe Ateniese
Johns Hopkins University – Baltimore, US
- Johannes Blömer
Universität Paderborn, DE
- Nikita Borisov
Univ. of Illinois – Urbana, US
- Xavier Boyen
Prime Cryptography – Palo Alto, US
- Jan Camenisch
IBM Research – Zürich, CH
- Claude Castelluccia
INRIA Rhône-Alpes, FR
- Bruno Crispo
University of Trento – Povo, IT
- George Danezis
Microsoft Research UK – Cambridge, GB
- Emiliano De Cristofaro
PARC – Palo Alto, US
- Claudia Diaz
K.U. Leuven, BE
- Roger Dingledine
The Tor Project, US
- Yevgeniy Dodis
New York University, US
- Maria Dubovitskaya
IBM Research – Zürich, CH
- Marc Fischlin
TU Darmstadt, DE
- Cédric Fournet
Microsoft Research UK – Cambridge, GB
- Ian Goldberg
University of Waterloo, CA
- Dennis Hofheinz
KIT – Karlsruhe Institute of Technology, DE
- Jean Pierre Hubaux
EPFL – Lausanne, CH
- Aaron M. Johnson
Naval Res. – Washington, US
- Stefan Katzenbeisser
TU Darmstadt, DE
- Dogan Kesdogan
Universität Siegen, DE
- Markulf Kohlweiss
Microsoft Research UK – Cambridge, GB
- Anja Lehmann
IBM Research – Zürich, CH
- Anna Lysyanskaya
Brown Univ. – Providence, US
- Mark Manulis
University of Surrey, GB
- Gregory Neven
IBM Research – Zürich, CH
- Bertram Poettering
RHUL – London, GB
- Bartosz Przydatek
Google Switzerland – Zürich, CH
- Kai Rannenberg
Goethe-Universität Frankfurt am Main, DE
- Jean-Pierre Seifert
TU Berlin, DE
- Vitaly Shmatikov
University of Texas at Austin, US
- Radu Sion
Stony Brook University, US
- Douglas Stebila
University of Technology Brisbane, AU
- Gene Tsudik
Univ. of California – Irvine, US
- Markus Ullmann
BSI – Bonn, DE
- Susanne Wetzel
Stevens Inst. of Technology, US
- Douglas Wikström
KTH Stockholm, SE
- Rebecca Wright
Rutgers Univ. – Piscataway, US