08491 Abstracts Collection
Theoretical Foundations of Practical Information Security
— Dagstuhl Seminar —

Ran Canetti[1], Shafi Goldwasser[2], Günter Müller[3] and Rainer Steinwandt[4]

[1] Tel Aviv University
canetti@post.tau.ac.il
[2] MIT - Cambridge, USA
shafi@theory.lcs.mit.edu
[3] Universität Freiburg, D
[4] Florida Atlantic University, USA
rsteinwa@fau.edu

Abstract. From 30.11. to 05.12.2008, the Dagstuhl Seminar 08491 "Theoretical Foundations of Practical Information Security " was held in Schloss Dagstuhl – Leibniz Center for Informatics. During the seminar, several participants presented their current research, and ongoing work and open problems were discussed. Abstracts of the presentations given during the seminar as well as abstracts of seminar results and ideas are put together in this paper. The first section describes the seminar topics and goals in general. Links to extended abstracts or full papers are provided, if available.

Keywords. Organic computing, self-organisation, design, adaptivity

08491 Executive Summary – Theoretical Foundations of Practical Information Security

Designing, building, and operating secure information processing systems is a complex task, and the only scientific way to address the diverse challenges arising throughout the life-cycle of security criticial systems is to consolidate and increase the knowledge of the theoretical foundations of practical security problems. To this aim, the mutual exchange of ideas across individual security research communities can be extraordinary beneficial. Accordingly, the motivation of this Dagstuhl seminar was the integration of different research areas with the common goal of providing an integral theoretical basis that is needed for the design of secure information processing systems.

Dagstuhl Seminar Proceedings 08491
Theoretical Foundations of Practical Information Security
http://drops.dagstuhl.de/opus/volltexte/2009/1894

Keywords: Organic computing, self-organisation, design, adaptivity

Joint work of: Canetti, Ran; Goldwasser, Shafi; Müller, Günter; Steinwandt, Rainer

Extended Abstract: http://drops.dagstuhl.de/opus/volltexte/2009/1893

Theoretical Foundations of Privacy Notions

Jens-Matthias Bohli (NEC Laboratories Europe - Heidelberg, DE)

We present a hierarchy of privacy notions that covers multiple anonymity and unlinkability variants. The underlying definitions, which are based on the idea of indistinguishability between two worlds, provide new insights into the relation between, and the fundamental structure of, different privacy notions. We apply the definitions to a number of privacy-preserving systems, namely group signatures, voting systems, and anonymous communication systems, and show how they relate to existing definitions.

Keywords: Privacy, group signatures, anonymous channels

Joint work of: Bohli, Jens-Matthias; Pashalidis, Andreas

Weak Verifiable Pseudorandom Functions

Zvika Brakerski (Weizmann Institute - Rehovot, IL)

Verifiable random functions (VRFs), introduced by Micali, Rabin and Vadhan, are pseudorandom functions in which the owner of the seed produces a public-key that constitutes a commitment to all values of the function. It can then produce, for any input x, a proof that the function has been evaluated correctly on x, preserving pseudorandomness for all other inputs. No public-key (even a falsely generated one) allows for proving more than one value per input (compare to signature schemes where soundness is only required for legal public-keys).

VRFs are both a natural and a useful primitive, and previous works have suggested a variety of constructions and applications. Still, there are many open questions in the study of VRFs, especially their relation to more widely studied cryptographic primitives and constructing them from a wide variety of cryptographic assumptions.

In this work we define a natural relaxation of VRFs that we call weak verifiable random functions, where pseudorandomness is required to hold only for randomly selected inputs. We conduct a study of weak VRFs, focusing on applications, constructions, and their relationship to other cryptographic primitives. We show:

- **Constructions.** We present constructions of weak VRFs based on a variety of assumptions, including general assumptions such as (enhanced) trapdoor permutations, as well as constructions based on specific number-theoretic assumptions such as the Diffie-Hellman assumption in bilinear groups.
- **Separations.** Verifiable random functions (both weak and standard) cannot be constructed from one-way permutations in a black-box manner. This constitutes the first result separating (standard) VRFs from any cryptographic primitive. While pseudorandom functions (without verifiability) and verifiable functions (or signature schemes) can each be constructed from one-way functions (using black-box reductions), we show that verifiable random functions (both weak and standard) cannot be constructed from one-way permutations in a black-box manner. This constitutes the first result separating (standard) VRFs from any cryptographic primitive.
- **Applications.** Weak VRFs capture the essence of constructing non-interactive zero-knowledge proofs for all NP languages.

Joint work of: Brakerski, Zvika; Goldwasser, Shafi; Rothblum, Guy; Vaikuntanathan, Vinod

Reduction of the Threshold Gap in Algebraic Geometric Secret Sharing

Ronald Cramer (CWI - Amsterdam, NL)

With Hao Chen (Shanghai) I showed that "quasi-threshold" linear secret sharing schemes with strong multiplication and corruption tolerance arbitrarily close to the limit 1/3n, exist over *fixed* base fields, using asymptotically good towers of algebraic function fields (CRYPTO 2006).

As a result, secure computation can be performed over fixed base fields, instead of over a field whose cardinality is linear in the size of the network. In certain scenarios this gives important savings in the amount of communication.

In recent joint work with Hao Chen, Chaoping Xing (Singapore) and Nacho Cascudo Pueyo (Oviedo), I showed that the (unavoidable) gap between privacy and reconstruction in such schemes can be reduced and that in fact such schemes exists over any finite field (not just the ones whose cardinality is a square). This is all by more advanced methods from the theory of function fields (compared to our earlier basic construction). We also improved a lower bound due to Karchmer and Wigderson (1993), and showed that if the gap does not grow fast enough, any linear secret sharing scheme respecting the gap will have vanishing information rate.

One-Round Authenticated Key Agreement from Weak Secrets

Yevgeniy Dodis (Courant Institute - New York, US)

We study the question of information-theoretically secure authenticated key agreement from weak secrets. In this setting, Alice and Bob share a n-bit secret W, which might *not* be uniformly random but the adversary has at least k bits of uncertainty about it (formalized using conditional min-entropy).

Alice and Bob wish to use W to agree on a nearly uniform secret key R, over a public channel controlled by an *active* adversary Eve. We show that non-interactive (single-message) protocols do not work when $k <= n/2$, and require poor parameters even when $n/2 < k << n$.

On the other hand, for arbitrary values of k, we design a communication efficient two-message (i.e, **one-round**!) protocol extracting nearly k random bits. This dramatically improves the only previously known protocol of Renner and Wolf [RW03], which required O(s) rounds where s is the security parameter. Our solution takes a new approach by studying and constructing **non-malleable seeded randomness extractors ** — if an attacker sees a random seed X and comes up with an arbitrarily related seed X', then we bound the relationship between $R = Ext(W; X)$ and $R' = Ext(W; X')$.

We also extend our one-round key agreement protocol to the "fuzzy" setting, where Alice and Bob share "close" (but not equal) secrets W_A and W_B, and to the Bounded Retrieval Model (BRM) where the size of the secret W is huge.

Keywords: Authenticated key agreement, information-theoretic security, non-malleable extractors

Accessible Entropy

Iftach Haitner (Microsoft Research New England - Cambridge, US)

We put forth a new computational notion of entropy, which measures the (in)-feasibility of sampling high entropy strings that are consistent with a given protocol. Specifically, we say that the i'th round of a protocol (A,B) has *accessible entropy* at most k, if no polynomial-time strategy A* can generate messages for A such that the entropy of its message in the iŠth round has entropy greater than k when conditioned both on prior messages of the protocol and on prior coin tosses of A*.

As applications of this notion, we

- Give a much simpler and more efficient construction of statistically hiding commitment schemes from arbitrary one-way functions, and
- Prove that constant-round statistically hiding commitments are necessary for constructing constant-round zero-knowledge proof systems for NP that remain secure under parallel composition (assuming the existence of one-way functions).

Joint work with Omer Reingold, Salil Vadhan and Hoeteck Wee.

Keywords: Computational complexity, cryptography, commitment schemes, interactive hashing, zero knowledge, one-way functions

Joint work of: Haitner, Iftach; Reingold, Omer; Vadhan, Salil; Wee, Hoeteck

Practical chosen-ciphertext secure encryption from factoring

Dennis Hofheinz (CWI - Amsterdam, NL)

We propose a new public-key encryption scheme. The scheme is practical and its security against chosen-ciphertext attacks can be reduced in the standard model to the assumption that factoring is intractable.

Keywords: Public-key encryption, chosen-ciphertext security, factoring

Joint work of: Hofheinz, Dennis; Kiltz, Eike

Partial Fairness in Secure Two-Party Computation

Jonathan Katz (University of Maryland - College Park, US)

Complete fairness is impossible to achieve, in general, in secure two-party computation. In light of this, various techniques for obtaining *partial* fairness in this setting have been suggested. We explore the possibility of achieving partial fairness with respect to a strong, simulation-based definition of security within the standard real/ideal world paradigm. We show feasibility with respect to this definition for randomized functionalities where each player may possibly receive a different output, as long as at least one of the domains or ranges of the functionality are polynomial in size. When one of the domains is polynomial size, our protocol is also secure-with-abort.

In contrast to much of the earlier work on partial fairness, we rely on standard assumptions only (namely, enhanced trapdoor permutations).

We also show that our results are optimal. Specifically, we show a boolean function defined on domains of super-polynomial size for which it is impossible to achieve both partial fairness and security with abort, and show a function whose domains and ranges all have super-polynomial size for which partial fairness is impossible altogether.

Keywords: Secure computation, fairness

Joint work of: Gordon, Dov; Katz, Jonathan

Sound and Fine-grain Specification of Ideal Functionalities

Aggelos Kiayias (University of Connecticut - Storrs, US)

Nowadays it is widely accepted to formulate the security of a protocol carrying out a given task via the "trusted-party paradigm", where the protocol execution is compared with an ideal process where the outputs are computed by a trusted party that sees all the inputs. A protocol is said to securely carry out a given task if running the protocol with a realistic adversary amounts to "emulating" the ideal process with the appropriate trusted party. In the Universal Composability (UC) framework the program run by the trusted party is called an ideal functionality. While this simulation-based security formulation provides strong security guarantees, its usefulness is contingent on the properties and correct specification of the ideal functionality, which, as demonstrated in recent years by the coexistence of complex, multiple functionalities for the same task as well as by their "unstable" nature, does not seem to be an easy task.

In this paper we address this problem, by introducing a general methodology for the sound specification of ideal functionalities. First, we introduce the class of canonical ideal functionalities for a cryptographic task, which unifies the syntactic specification of a large class of cryptographic tasks under the same basic template functionality. Furthermore, this representation enables the isolation of the individual properties of a cryptographic task as separate members of the corresponding class. By endowing the class of canonical functionalities with an algebraic structure we are able to combine basic functionalities to a single final canonical functionality for a given task. Effectively, this puts forth a bottom-up approach for the specification of ideal functionalities: first one defines a set of basic constituent functionalities for the task at hand, and then combines them into a single ideal functionality taking advantage of the algebraic structure.

In our framework, the constituent functionalities of a task can be derived either directly or, following a translation strategy we introduce, from existing game-based definitions; such definitions have in many cases captured desired individual properties of cryptographic tasks, albeit in less adversarial settings than universal composition. Our translation methodology entails a sequence of steps that derive a corresponding canonical functionality given a game-based definition. In this way, we obtain a well-defined mapping of game-based security properties to their corresponding UC counterparts.

Finally, we demonstrate the power of our approach by applying our methodology to a variety of basic cryptographic tasks, including commitments, digital signatures, zero-knowledge proofs, and oblivious transfer.

While in some cases our derived canonical functionalities are equivalent to existing formulations, thus attesting to the validity of our approach, in others they differ, enabling us to "debug" previous definitions and pinpoint their shortcomings.

Keywords: Security definitions, universal composability, cryptographic protocols, lattices and partial orders

Joint work of: Garay, Juan; Kiayias, Aggelos; Zhou, Hong-Sheng

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1891

Joint State Theorems for Public-Key Encryption and Digitial Signature Functionalities with Local Computation

Ralf Küsters (Universität Trier, DE)

Composition theorems in simulation-based approaches allow to build complex protocols from sub-protocols in a modular way. However, as first pointed out and studied by Canetti and Rabin, this modular approach often leads to impractical implementations. For example, when using a functionality for digital signatures within a more complex protocol, parties have to generate new verification and signing keys for every session of the protocol. This motivates to generalize composition theorems to so-called joint state theorems, where different copies of a functionality may share some state, e.g., the same verification and signing keys.

In this work, we present a joint state theorem which is more general than the original theorem of Canetti and Rabin, for which several problems and limitations are pointed out. We apply our theorem to obtain joint state realizations for three functionalities: public-key encryption, replayable public-key encryption, and digital signatures. Unlike most other formulations, our functionalities model that ciphertexts and signatures are computed locally, rather than being provided by the adversary. To obtain the joint state realizations, the functionalities have to be designed carefully. Other formulations are shown to be unsuitable. Our work is based on a recently proposed, rigorous model for simulation-based security by Küsters, called the IITM model. Our definitions and results demonstrate the expressivity and simplicity of this model. For example, unlike Canetti's UC model, in the IITM model no explicit joint state operator needs to be defined and the joint state theorem follows immediately from the composition theorem in the IITM model.

Joint work with Max Tuengerthal (CSF 2008)

Keywords: Simulation-based security, joint state theorem

Joint work of: Küsters, Ralf; Tuengerthal, Max

Simple, Black-Box Constructions of Adaptively Secure Protocols

Tal Malkin (Columbia University, US)

We present a new compiler for transforming an oblivious transfer (OT) protocol secure against an adaptive semi-honest adversary into one that is secure against an adaptive malicious adversary.

Our compiler achieves security in the universal composability framework, assuming access to an ideal commitment functionality. Our compiler improves over previous work achieving the same security guarantee in two ways: it uses black-box access to the underlying protocol and achieves a constant multiplicative overhead in the round complexity. As a corollary, we obtain the first constructions of adaptively secure protocols in the stand-alone model using black-box access to a low-level primitive.

Joint work of: Choi, Seung Geol; Dachman-Soled, Dana; Malkin, Tal; Wee, Hoeteck

Wireless Physical Layer Key Exchange

Jörn Müller-Quade (Universität Karlsruhe, DE)

Secure wireless key on the physical layer allows a key exchange independent of computational assumptions. The precise conditions, however, for such a key exchange to be secure remain unclear. In this work we claim that a security analysis of such protocols must be done on a rigorous physical layer, because the abstract stochastic models used in communication theory might overestimate the level of noise present. This is no problem for error correction applications, but for security protocols this noise is a resource and it should rather be conservatively underestimated to obtain reliable security gurantees.

In the concrete example of a key exchange based on multipath interference and the reciprocity of the wireless channel we show that (in a simplified model) each single signal received by an eavesdroper is uncorrelated to the key obtained by the legitimite players (Alice and Bob). This uncorrelatedness is (implicitly) used as an argument fort he security of the key exchange. However, as soon as Eve receives two signals (as it ist he case for reciprocity based key exchange) the eavesdropper can reconstruct the key. Different uses of a channel may hence not be modelled by uncorrelated stochastic processes. Furthermore we present a side channel attack on reciprocity based key exchange which was neglected so far.

A novel reciprocity based key exchange protocol is introduced which seems to cope with the problems presented in this work, but a proper model of security to prove such claims remains an important task for future work.

Keywords: Wireless, Reciprocity, Key Exchange

Joint work of: Almeida, Antonio, Döttling, Nico, Gabel, Matthias, Müller-Quade, Jörn

How Fair Can a Coin Toss Be?

Moni Naor (Weizmann Institute - Rehovot, IL)

Coin-flipping protocols allow mutually distrustful parties to generate a common unbiased random bit, guaranteeing that even if one of the parties is malicious, it cannot significantly bias the output of the honest party.

A classical result by Cleve from STOC 1986 showed that without simultaneous broadcast, for any two-party coin-flipping protocol with r rounds, there exists an efficient adversary that can bias the output of the honest party by Omega(1/r). However, the best previously known protocol only guarantees O(1/sqrt(r)) bias and Cleve and Impagliazzo have shown that this is optimal in the fail-stop model.

We establish the optimal tradeoff between the round complexity and the maximal bias of two-party coin-flipping protocols. Under standard assumptions, we show that Cleve's lower bound is tight: we construct an r-round protocol with bias O(1/r). We make use of recent progress by Gordon, Hazay, Katz and Lindell regarding fair protocols (STOC 2008).

Joint work with Tal Moran and Gil Segev

Joint work of: Tal Moran, Moni Naor and Gil Segev

Full Paper:
http://www.wisdom.weizmann.ac.il/%7Enaor/PAPERS/optimal_coin.pdf

Game Theory with Costly Computation

Rafael Pass (Cornell University - Ithaca, US)

We develop a general game-theoretic framework for reasoning about strategic agents performing possibly costly computation. In this framework, many traditional game-theoretic results (such as the existence of a Nash equilibrium) no longer hold. Nevertheless, we can use the framework to provide psychologically appealing explanations to observed behavior in well-studied games (such as finitely repeated prisoner's dilemma and rock-paper-scissors). Furthermore, we provide natural conditions on games sufficient to guarantee that equilibria exist.

As an application of this framework, we provide a game-theoretic definition of protocol security. We show that a special case of this notion is equivalent to a variant of the traditional cryptographic definition of protocol security; this results shows that, when taking computation into account, the two approaches used for dealing with "deviating" players in two different communities—Nash equilibrium in game theory, and zero-knowledge "simulation" in cryptography—are intimately related.

Joint work with Joe Halpern

Public-Key Cryptosystems from the Worst-Case Shortest Vector Problem

Chris Peikert (SRI - Menlo Park, US)

We construct public-key cryptosystems that are secure assuming the *worst-case* hardness of approximating the shortest vector problem on lattices. Prior cryptosystems with worst-case connections (e.g., the Ajtai-Dwork system) were based either on a *special case* of the shortest vector problem, or on the conjectured hardness of lattice problems for *quantum* algorithms.

Our main technical innovation is a reduction from certain variants of the shortest vector problem to corresponding versions of the "learning with errors" (LWE) problem; previously, only a quantum reduction of this kind was known. In addition, we construct new cryptosystems based on LWE, including a very natural chosen ciphertext-secure system that has a much simpler description and tighter underlying worst-case approximation factor than prior constructions.

Keywords: Lattice-based cryptography, learning with errors, quantum computation

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1892

Modeling Computational Security in Long-Lived Systems

Olivier Pereira (University of Louvain, BE)

For many cryptographic protocols, security relies on the assumption that adversarial entities have limited computational power.

This type of security degrades progressively over the lifetime of a protocol. However, some cryptographic services, such as timestamping services or digital archives, are *long-lived* in nature; they are expected to be secure and operational for a very long time (i.e. super-polynomial). In such cases, security cannot be guaranteed in the traditional sense: a computationally secure protocol may become insecure if the attacker has a super-polynomial number of interactions with the protocol. This paper proposes a new paradigm for the analysis of long-lived security protocols. We allow entities to be active for a potentially unbounded amount of real time, provided they perform only a polynomial amount of work *per unit of real time*. Moreover, the space used by these entities is allocated dynamically and must be polynomially bounded. We propose a new notion of *long-term implementation*, which is an adaptation of computational indistinguishability to the long-lived setting. We show that long-term implementation is preserved under polynomial parallel composition and exponential sequential composition.

We illustrate the use of this new paradigm by analyzing some security properties of the long-lived timestamping protocol of Haber and Kamat.

Keywords: Long lived security, universally composable security

Joint work of: Canetti, Ran; Cheung, Ling; Kaynar, Dilsun; Lynch, Nancy; Pereira, Olivier

Full Paper: http://drops.dagstuhl.de/opus/volltexte/2009/1890

Full Paper:
http://eprint.iacr.org/2008/492

Leakage-Resilient Cryptography: Theoretical Foundations of Side-Channel Security

Krzysztof Pietrzak (CWI - Amsterdam, NL)

Below is the abstract of a joint paper with Stefan Dziembowski (FOCS'08) on which much of the talk is based.

We construct a stream-cipher SC whose *implementation* is secure even if a bounded amount of arbitrary (adversarially chosen) information on the internal state of SC is leaked during computation. This captures *all* possible side-channel attacks on SC where the amount of information leaked in a given period is bounded, but overall can be arbitrary large. The only other assumption we make on the *implementation* of $\mathcal{S}$ is that only data that is accessed during computation leaks information.

The stream-cipher SC generates its output in chunks $K_1, K_2, \ldots$ and arbitrary but bounded information leakage is modeled by allowing the adversary to adaptively chose a function $f_\ell : \{0,1\}^* \to \{0,1\}^\lambda$ before K_ℓ is computed, she then gets $f_\ell(\tau_\ell)$ where τ_ℓ is the internal state of $\mathcal{S}$ that is accessed during the computation of K_ℓ.

One notion of security we prove for $\mathcal{S}$ is that K_ℓ is indistinguishable from random when given $K_1, \ldots, K_{\ell-1}$, $f_1(\tau_1), \ldots, f_{\ell-1}(\tau_{\ell-1})$ and also the complete internal state of $\mathcal{S}$ after K_ℓ has been computed (i.e. $\mathcal{S}$ is forward-secure).

The construction is based on alternating extraction (used in the intrusion-resilient secret-sharing scheme from FOCS'07). We move this concept to the computational setting by proving a lemma that states that the output of any PRG has high HILL pseudoentropy (i.e. is indistinguishable from some distribution with high min-entropy) even if arbitrary information about the seed is leaked. The amount of leakage λ that we can tolerate in each step depends on the strength of the underlying PRG, it is at least logarithmic, but can be as large as a constant fraction of the internal state of $\mathcal{S}$ if the PRG is exponentially hard.

Keywords: Side-channel, stream-cipher, leakage-resilience, extractors

Joint work of: Pietrzak, Krzysztof; Dziembowski, Stefan

Founding Cryptography on Oblivious Transfer – Efficiently

Manoj Prabhakaran (University of Illinois - Urbana, US)

We present a simple and efficient compiler for transforming secure multi-party computation (MPC) protocols that enjoy security only with an honest majority into MPC protocols that guarantee security with no honest majority, in the oblivious-transfer (OT) hybrid model. Our technique works by combining a secure protocol in the honest majority setting with a protocol achieving only security against *semi-honest* parties in the setting of no honest majority.

Applying our compiler to variants of protocols from the literature, we get several applications for secure two-party computation and for MPC with no honest majority. These include:

- **Constant-rate two-party computation in the OT-hybrid model.** We obtain a statistically UC-secure two-party protocol in the OT-hybrid model that can evaluate a general circuit C of size s and depth d with a total communication complexity of $O(s)+(k, d, \log s)$ and $O(d)$ rounds. The above result generalizes to a constant number of parties.
- **Extending OTs in the malicious model.** We obtain a computationally efficient protocol for generating many string OTs from few string OTs with only a *constant amortized communication overhead* compared to the total length of the string OTs.
- **Black-box constructions for constant-round MPC with no honest majority.** We obtain general computationally UC-secure MPC protocols in the OT-hybrid model that use only a constant number of rounds, and only make a *black-box* access to a pseudorandom generator. This gives the first constant-round protocols for three or more parties that only make a black-box use of cryptographic primitives (and avoid expensive zero-knowledge proofs).

Keywords: Multi-party computation, Oblivious transfer, UC security

See also: @inproceedingsDBLP:conf/crypto/IshaiPS08, author = Yuval Ishai and Manoj Prabhakaran and Amit Sahai, title = Founding Cryptography on Oblivious Transfer - Efficiently, booktitle = CRYPTO, year = 2008, pages = 572-591, ee = http://dx.doi.org/10.1007/978-3-540-85174-5_32,
crossref = DBLP:conf/crypto/2008, bibsource DBLP, http://dblp.uni-trier.de

Fairness with an honest minority and a rational majority

Alon Rosen (The Interdisciplinary Center - Herzliya, IL)

I will present a simple protocol for secret reconstruction in any threshold secret sharing scheme, and argue that it is fair when executed with many rational parties together with a small minority of honest parties.

That is, all parties will learn the secret with high probability when the honest parties follow the protocol and the rational parties act in their own self-interest. The protocol only requires a standard (synchronous) broadcast channel, and tolerates arbitrary deviations (including early stopping and incorrectly computed messages).

Previous protocols for this problem in the cryptographic or economic models have either required an honest majority, used strong communication channels that enable simultaneous exchange of information, or settled for approximate notions of security/equilibria.

Joint work with Shien Jin Ong, David Parkes and Salil Vadhan.

When and How Can Data be Efficiently Released with Privacy?

Guy Rothblum (MIT - Cambridge, US)

We consider private data analysis in the setting in which a trusted and trustworthy curator, having obtained a data set containing sensitive information, releases to the public a "sanitization" of the data set. The goal is for the sanitization to both protect the privacy of the individual contributors of data and offer aggregate statistical utility to a data analyst. We focus on the case where the process is non-interactive; once the sanitization has been released the original data and the curator play no further role.

Blum et al. [STOC '08] showed a remarkable result: for any collection of counting predicate queries, the exponential mechanism of McSherry and Talwar [FOCS '07] can be used to (inefficiently) generate a synthetic dataset that maintains approximately correct fractional counts for all of the queries, while ensuring a strong privacy guarantee, known as differential privacy.

In this work we investigate the computational complexity of such non-interactive privacy mechanisms, mapping the boundary between feasibility and infeasibility. We show:

1. For any query set C and data universe X there is an efficient sanitizer that outputs a synthetic dataset and runs in time poly(|C|,|X|). The sanitizer works even when the size of the original dataset is sub-polynomial in |C|.

2. Under standard cryptographic assumptions, there is no sanitizer that outputs synthetic datasets and works for general classes of counting queries and runs in time sub-polynomial in |C| or |X|. In particular, this is a case where the exponential mechanism cannot, in general, be implemented efficiently.

3. Turning to the potentially easier problem of privately generating an arbitrary data structure (not necessarily synthetic data), there is a tight connection between hardness of sanitization and the existence of traitor tracing schemes.

Joint work of: Dwork, Cynthia; Naor, Moni; Reingold, Omer; Rothblum, Guy; Vadhan, Salil

Chosen-Ciphertext Security via Correlated Products

Gil Segev (Weizmann Institute - Rehovot, IL)

We initiate the study of one-wayness under "correlated products". We are interested in identifying necessary and sufficient conditions for a function f and a distribution on inputs $(x_1, ..., x_k)$, so that the function $(f(x_1), ..., f(x_k))$ is one-way. The main motivation of this study is the construction of public-key encryption schemes that are secure against chosen-ciphertext attacks (CCA). We show that any collection of injective trapdoor functions that is secure under a very natural correlated product can be used to construct a CCA-secure public-key encryption scheme. The construction is simple, black-box, and admits a direct proof of security.

We provide evidence that security under correlated products is achievable by demonstrating that lossy trapdoor functions (Peikert and Waters, STOC '08) yield injective trapdoor functions that are secure under the above mentioned correlated product. Although we eventually base security under correlated products on existing constructions of lossy trapdoor functions, we argue that the former notion is potentially weaker as a general assumption. Specifically, there is no fully-black-box construction of lossy trapdoor functions from trapdoor functions that are secure under correlated products.

Joint work of: Rosen, Alon; Segev, Gil

Abstractions for Cryptographically Faithful Proofs of Security Protocols

Christoph Sprenger (ETH Zürich, CH)

In this talk, I will present a formal theory for cryptographically-sound theorem proving. Our starting point is the Backes-Pfitzmann-Waidner (BPW) model, which is a symbolic protocol model that is cryptographically sound in the sense of blackbox reactive simulatability. To achieve cryptographic soundness, this model is substantially more complex than standard symbolic models and the main challenge in formalizing and using this model is overcoming this complexity. We have constructed a series of cryptographically-sound abstractions of the original BPW model that bring it much closer to standard Dolev-Yao style models. I will present a case study showing that our abstractions enable proofs of size and complexity comparable to those based on Paulson's much simpler inductive method. Our entire development has been formalized in the theorem prover Isabelle/HOL.

Based on joint work with Michael Backes, David Basin, Birgit Pfitzmann and Michael Waidner.

Keywords: Security protocols, cryptographic soundness, reactive simulatability, theorem proving, higher-order logic

A Functionality for Symmetric Encryption in Simulation-based Security

Max Tuengerthal (Universität Trier, DE)

For most basic cryptographic tasks, such as public-key encryption, digital signatures, authentication, key exchange, and many other more sophisticated tasks, ideal functionalities in the simulation-based security approach have been formulated, along with their realizations.

Surprisingly, however, no such functionality exists for symmetric encryption, except for a more abstract Dolev-Yao style functionality.

In this work, we fill this gap. We propose two functionalities for symmetric encryption, an unauthenticated and an authenticated version, and show that they can be implemented based on standard cryptographic assumptions for symmetric encryption schemes, namely IND-CCA security and authenticated encryption, respectively. We also illustrate the usefulness of our functionalities in applications, both in simulation-based and game-based security settings.

Keywords: Cryptographic protocols, simulation-based security

Joint work of: Kuesters, Ralf; Tuengerthal, Max

Security of public key encryption under key dependent messages

Dominique Unruh (Universität des Saarlandes, DE)

Commonly, security of encryption schemes does not give any guarantees if the plaintext may include or otherwise depend on the secret key. Encryption schemes that are secure in the presence of key dependent messages were so far only studied in the setting of private key cryptography. We explain why adapting the existing definitions to the public key setting is difficult and give the first security definition for public key encryption under key dependent messages. We show that the OAEP encryption scheme is secure under key dependent messages in the random oracle model.

Keywords: Encryption schemes, key dependent messages

Joint work of: Michael Backes; Markus Dürmuth; Unruh, Dominique

Full Paper:
http://www.infsec.cs.uni-sb.de/∼unruh/publications/backes08oaep.html

See also: Michael Backes, Markus Dürmuth, and Dominique Unruh. OAEP is Secure Under Key-dependent Messages. ASIACRYPT 2008

Efficient and Composable Oblivious Transfer

Vinod Vaikuntanathan (MIT - Cambridge, US)

We propose a simple and general framework for constructing oblivious transfer (OT) protocols that are efficient, universally composable, and generally realizable from a variety of cryptographic assumptions, such as the decisional Diffie-Hellman assumption, the Quadratic Residuosity assumption and worst-case complexity assumptions relating to lattices. Our OT protocols are round-optimal (one message each way) and efficient in the parties' communication and local computation.

One of our key technical contributions is a unified view of several encryption schemes in the literature that have what we call message-lossy public keys, whose defining property is that a cipher-text produced under such a key carries no information (even statistically) about the encrypted message.

Keywords: Oblivious Transfer, UC security

Joint work of: Peikert, Chris; Vaikuntanathan, Vinod; Waters, Brent

Full Paper:
http://www.mit.edu/~vinodv/papers/OT.pdf

Non-malleable Obfuscation

Mayank Varia (MIT - Cambridge, US)

Existing definitions of program obfuscation do not rule out malleability attacks, where an adversary that sees an obfuscated program is able to generate another (potentially obfuscated) program that is related to the original one in some way.

We formulate two natural flavors of non-malleability requirements for program obfuscation, and show that they are incomparable in general. We also construct non-malleable obfuscators of both flavors for some program families of interest. Some of our constructions are in the Random Oracle model, whereas another one is in the common reference string model. We also define the notion of verifiable obfuscation which is of independent interest.

Keywords: Obfuscation, non-malleable

Joint work of: Canetti, Ran; Varia, Mayank