Report from Dagstuhl Seminar 16371

Public-Key Cryptography

Edited by

Marc Fischlin[1], Alexander May[2], David Pointcheval[3], and Tal Rabin[4]

1   TU Darmstadt, DE, `marc.fischlin@cryptoplexity.de`
2   Ruhr-Universität Bochum, DE, `alex.may@ruhr-uni-bochum.de`
3   ENS – Paris, FR, `david.pointcheval@ens.fr`
4   IBM Thomas J. Research Center – Yorktown Heights, US, `talr@us.ibm.com`

---- **Abstract** --

This report documents the program and results of Dagstuhl seminar 16731 "Public-Key Cryptography" which took place September 11–16, 2016. The goal of the seminar was to bring together different subareas from public-key cryptography and to promote research among these areas.

1   Summary

Marc Fischlin

Cryptography has turned out to be an invaluable tool for protecting the confidentiality and integrity of digital data. At the same time, cryptography does not yet provide satisfying solutions to all practical scenarios and threats. To accomplish appropriate protection of the data, cryptography needs to address several challenges.

Cryptography has always been a prominent theme within the Dagstuhl Seminar series, with the first meeting about cryptography held in 1993, and subsequent seminars on this topic about every 5 years. In 2007 and 2012 a seminar for the subarea of "Symmetric Cryptography" has been added, inciting us to coin the seminar here "Public-Key Cryptography" for sake of distinction. The public-key branch has been held for the second time, after the first event in 2011.

The seminar brought together 27 scientists in the area of public-key cryptography, including three student researchers who were invited by Dagstuhl to pick a seminar to participate in. The participants came from all over the world, including countries like the US, Great Britain, Israel, France, or Japan. Among the affiliations, Germany lead the number with 9 participants, followed by the US and France with 6 each. The program contained 21 talks, each of 25 to 60 minutes, and a panel discussion about the uneasiness with the current state of our reviewing system, with a free afternoon on Wednesday for social activities and the afternoon on Thursday for collaborations. Before the seminar, we asked the participants to present very recent and ongoing work which, ideally, should not have been published or

accepted to publication yet. Most of the participants followed our suggestion and to a large extend the presentations covered topics which have not even been submitted at the time.

The topics of the talks represented the diversity of public-key cryptography. The goal of the seminar was to bring together three challenge areas in cryptography, namely, cryptanalysis and foundations (investigating and evaluating new primitives), optimization (making solutions more efficient), and deployment (designing real-world protocols). As envisioned, the seminar thus has a good mixture of talks from these areas. There were also suggestions to try to co-locate future events of the seminar with other security-related events at Dagstuhl to foster even broader interdisciplinary research. Discussions during and after the talks were lively. It seems as if the goal of stimulating collaborations among these areas has been met. The discussion about the reviewing system has led to some hands-on practices which could be deployed to improve the quality of reviews. This includes incentives such as"Best Reviewer Awards" and teaching students about proper reviewing.

2   Table of Contents

3 Overview of Talks

3.1 Diverse Vector Spaces and Zero-Knowledge

Fabrice Benhamouda (IBM Thomas J. Watson Research Center – Yorktown Heights, US)

Joint work of Michel Abdallah, Fabrice Benhamouda, David Pointcheval
Main reference M. Abdallah, F. Benhamouda, D. Pointcheval, "Disjunctions for Hash Proof Systems: New
Constructions and Applications," in Proc. of the 34th Annual Int'l Conf. on the Theory and
Applications of Cryptographic Techniques – Advances in Cryptology (EUROCRYPT'15) – Part II,
LNCS, Vol. 9057, pp. 69–100, Springer, 2015; pre-print available at IACR.
URL http://dx.doi.org/10.1007/978-3-662-46803-6_3
URL https://eprint.iacr.org/2014/483

We first present hash proof systems or smooth projective hash functions (SPHFs), which were introduced by Cramer and Shoup in 2002 to explain the construction of the Cramer-Shoup IND-CCA encryption scheme and which later found numerous other applications. We then introduce diverse vector spaces as a tool to construct SPHFs. Finally, we illustrate this tool on simple examples and show applications to zero-knowledge primitives.

3.2 What Else is Revealed by Order-Revealing Encryption?

David Cash (Rutgers University, US)

Joint work of F. Betül Durak, Thomas M. DuBuisson, David Cash
Main reference F. Betül Durak, Thomas M. DuBuisson, David Cash, "What Else is Revealed by Order-Revealing
Encryption?," in Proc. of the 2016 ACM SIGSAC Conf. on Computer and Communications
Security (CCS'16), pp. 1155–1166, ACM, 2016; pre-print available at IACR.
URL http://dx.doi.org/10.1145/2976749.2978379
URL http://eprint.iacr.org/2016/786

The security of order-revealing encryption (ORE) has been unclear since its invention. Dataset characteristics for which ORE is especially insecure have been identified, such as small message spaces and low-entropy distributions. On the other hand, properties like one-wayness on uniformly-distributed datasets have been proved for ORE constructions.

This work shows that more plaintext information can be extracted from ORE ciphertexts than was previously thought. We identify two issues: First, we show that when multiple columns of correlated data are encrypted with ORE, attacks can use the encrypted columns together to reveal more information than prior attacks could extract from the columns individually. Second, we apply known attacks, and develop new attacks, to show that the leakage of concrete ORE schemes on non-uniform data leads to more accurate plaintext recovery than is suggested by the security theorems which only dealt with uniform inputs.

3.3 Comparison between Subfield and Straightforward Attacks on NTRU

Pierre-Alain Fouque (University of Rennes, FR)

Recently in two independent papers, Albrecht, Bai and Ducas and Cheon, Jeong and Lee presented two very similar attacks, that allow to break NTRU with larger parameters and GGH Multinear Map without zero encodings. They proposed an algorithm for recovering the NTRU secret key given the public key which apply for large NTRU modulus, in particular to Fully Homomorphic Encryption schemes based on NTRU. Hopefully, these attacks do not endanger the security of the NTRUE NCRYPT scheme, but shed new light on the hardness of this problem. The basic idea of both attacks relies on decreasing the dimension of the NTRU lattice using the multiplication matrix by the norm (resp. trace) of the public key in some subfield instead of the public key itself. Since the dimension of the subfield is smaller, the dimension of the lattice decreases, and lattice reduction algorithm will perform better. Here, we revisit the attacks on NTRU and propose another variant that is simpler and outperforms both of these attacks in practice. It allows to break several concrete instances of YASHE, a NTRU-based FHE scheme, but it is not as efficient as the hybrid method of Howgrave-Graham on concrete parameters of NTRU. Instead of using the norm and trace, we propose to use the multiplication by the public key in some subring and show that this choice leads to better attacks. We can then show that for power of two cyclotomic fields, the time complexity is polynomial. Finally, we show that, under heuristics, straightforward lattice reduction is even more efficient, allowing to extend this result to fields without non-trivial subfields, such as NTRU Prime. We insist that the improvement on the analysis applies even for relatively small modulus ; though if the secret is sparse, it may not be the fastest attack. We also derive a tight estimation of security for (Ring-)LWE and NTRU assumptions. when $q = 2^{\Omega(\sqrt{n \log \log n})}$.

3.4 Advances in building Non-Malleable Commitments

Vipul Goyal (Microsoft Research India – Bangalore, IN)

A central challenge in the design of secure systems is to defend against man-in-the-middle attacks, where an adversary can arbitrarily tamper with the messages exchanged by two parties over a communication channel. Starting with the early nineties, an important research goal in cryptography has been to build "non malleable" cryptographic protocols that are resilient to such attacks.

A very basic non-malleable primitive which is widely used in cryptography is what is known as non-malleable commitment schemes. In this talk, I will describe a recent result which constructs non-malleable commitments in the minimal number of rounds (and almost minimal complexity assumptions). In some sense, this culminates a two-decade long research quest of getting non-malleable commitments in the minimal number of rounds.

3.5 Fair Coin Flipping: Tighter Analysis and the Many-Party Case

Iftach Haitner (Tel Aviv University, IL)

In a multi-party fair coin-flipping protocol, the parties output a common (close to) unbiased bit, even when some corrupted parties try to bias the output. In this work we focus on the case of dishonest majority, i.e. at least half of the parties can be corrupted. Cleve (STOC 1986) has shown that in any m-round coin-flipping protocol the corrupted parties can bias the honest parties' common output bit by $1/m$. For more than two decades the best known coin-flipping protocols against dishonest majority was the protocol of Awerbuch, Blum, Chor, Goldwasser, and Micali [Manuscript 85], who presented a t-party, m-round protocol of bias $t/\sqrt{m}$. This was changed by the breakthrough result of Moran, Naor and Segev (TCC 2009), who constructed an m-round, 2-party coin-flipping protocol with optimal bias of $1/m$. Recently, Haitner and Tsafadia (STOC 14) constructed an m-round, three-party coin-flipping protocol with bias $O(log^3(m)/m)$. Still for the case of more than three parties, the best known protocol remains the $\Theta(t/\sqrt{m})$-bias protocol of Awerbuch et al.

We make a step towards eliminating the above gap, presenting a t-party, m-round coin-flipping protocol, with bias $O(\frac{t*2^t*\sqrt{\log m}}{m^{1/2+1/(2^{t-1}-2)}})$. This improves upon the $\Theta(t/\sqrt{m})$-bias protocol of Awerbuch et al. for any $t < 1/2 * log(log(m))$, and in particular for $t \in O(1)$, this yields an $1/m^{1/2+\Theta(1)}$-bias protocol. For the three-party case, this yields an $O(\sqrt{logm}/m)$-bias protocol, improving over the $O(log^3m/m)$-bias protocol of Haitner and Tsafadia. Our protocol generalizes that of Haitner and Tsafadia, by presenting an appropriate "defense protocols" for the remaining parties to interact in, in the case that some parties abort or caught cheating (Haitner and Tsafadia only presented a two-party defense protocol, which limits their final protocol to handle three parties).

We analyze our new protocols by presenting a new paradigm for analyzing fairness of coin-flipping protocols. We map the set of adversarial strategies that try to bias the honest parties outcome in the protocol to the set of the feasible solutions of a linear program. The gain each strategy achieves is the value of the corresponding solution. We then bound the optimal value of the linear program by constructing a feasible solution to its dual.

3.6 Kurosawa-Desmedt Meets Tight Security

Dennis Hofheinz (KIT – Karlsruher Institut für Technologie, DE)

At EUROCRYPT 2016, Gay et al. presented the first pairing-free public-key encryption (PKE) scheme with a tight security reduction to a standard assumption. Their scheme is competitive in efficiency with state-of-the art PKE schemes and has very compact ciphertexts (of three group elements), but suffers from a large public key (of about 200 group elements).

In this work, we present an improved pairing-free PKE scheme with a tight security reduction to the Decisional Diffie-Hellman assumption, small ciphertexts (of three group elements), *and* small public keys (of six group elements). Compared to the work of Gay et al.,

our scheme thus has a considerably smaller public key and comparable other characteristics, although our encryption and decryption algorithms are somewhat less efficient.

Technically, our scheme borrows ideas both from the work of Gay et al. and from a recent work of Hofheinz (eprint, 2016). The core technical novelty of our work is an efficient and compact designated-verifier proof system for an OR-like language. We show that adding such an OR-proof to the ciphertext of the state-of-the-art PKE scheme from Kurosawa and Desmedt enables a tight security reduction.

3.7   Schnorr Signatures in the Multi-User Setting

Eike Kiltz (Ruhr-Universität Bochum, DE)

Joint work of Eike Kiltz, Daniel Masny, Jiaxin Pan
Main reference E. Kiltz, D. Masny, J. Pan, "Optimal Security Proofs for Signatures from Identification Schemes,"
         in Proc. of the 36th Annual Int'l Cryptology Conf. – Advances in Cryptology (CRYPTO'16) –
         Part II, LNCS, Vol. 9815, pp. 33–61, Springer, 2016.
    URL http://dx.doi.org/10.1007/978-3-662-53008-5_2

A theorem by Galbraith, Malone-Lee, and Smart (GMLS) from 2002 showed that, for Schnorr signatures, single-user security tightly implies multi-user security. Recently, Bernstein pointed to an error in the above theorem and promoted a key-prefixing variant of Schnorr signatures for which he proved a tight implication from single to multi-user security. Even worse, he identified an "apparently insurmountable obstacle to the claimed [GMLS] theorem". This paper shows that, without key prefixing, single-user security of Schnorr signatures tightly implies multi-user security of the same scheme. Our result has slightly stronger requirements than the GLML theorem: we either require the random oracle model or strong single user security of Schnorr signatures.

3.8   Computational Arithmetic Secret Sharing

Alexander Koch (KIT – Karlsruher Institut für Technologie, DE)

Secret sharing schemes allow for sharing a secret message so that it can be correctly reconstructed in the presence of enough of its shares, but with the property that nothing can be learned about its content if too few of the shares have been obtained. Homomorphic schemes exhibit the additional property that it is possible to calculate on the shares to obtain a share of the sums and products of secrets, yielding a plethora of applications including secure multiparty computation (MPC). To reduce the size of the generated shares in a secret sharing scheme, "computational" variants have been developed which guarantee secrecy for illegitimate access to the secret only against computationally restricted adversaries. While these schemes are much more size-efficient, they usually have the disadvantage of not being homomorphic. We give the first computational secret sharing scheme on the basis of multi-key fully homomorphic encryption, that combines the advantages of both worlds.

3.9 Practical LPN Cryptanalysis

Alexander May (Ruhr-Universität Bochum, DE)

We present memory-efficient algorithms for LPN, both classically and quantumly. We also show first experiments for solving LPN instances up to dimension 250 with error parameter 1/8.

3.10 Concurrently Composable Security With Shielded Super-polynomial Simulators

Jörn Müller-Quade (KIT – Karlsruher Institut für Technologie, DE)

We propose a new framework for concurrently composable security that relaxes the security notion of UC security. As in previous frameworks, our notion is based on the idea of providing the simulator with super-polynomial resources. However, in our new framework simulators are only given *restricted access* to the results computed in super-polynomial time. This is done by modeling the super-polynomial resource as a stateful oracle that may directly interact with a functionality without the simulator seeing the communication. We call these oracles "shielded oracles".

Our notion is fully compatible with the UC framework, i.e., protocols proven secure in the UC framework remain secure in our framework. Furthermore, our notion lies strictly between SPS and Angel-based security, while being closed under protocol composition.

Shielding away super-polynomial resources allows us to apply new proof techniques where we can replace super-polynomial entities by indistinguishable polynomially bounded entities. This allows us to construct secure protocols in the plain model using weaker primitives than in previous composable frameworks involving simulators with super-poly resources. In particular, we only use non-adaptive-CCA-secure commitments as a building block in our constructions. As a feasibility result, we present a constant-round general MPC protocol in the plain model based on standard assumptions that is secure in our framework.

3.11 Lattice Enumeration Revisited

Phong Q. Nguyen (Inria and CNRS/JFLI, FR, and University of Tokyo, JP)

Lattice enumeration is arguably the simplest method to solve exact lattice problems. Though it does not have the best asymptotical time complexity, it has been used in the largest lattice records, notably NTRU challenges, Darmstadt's lattice challenges and SVP challenges. In this talk, we revisit lattice enumeration with pruning techniques.

3.12 Overcoming Hellman's Time/Memory Trade Offs with Applications to Proofs of Space

Krzysztof Pietrzak (IST Austria – Klosterneuburg, AT)

Hellman showed that any permutation over a domain of size N can be inverted in time T by an algorithm whose description is of size S for any S, T which satisfy $N < O(S \cdot T)$ (e.g. $S = T \approx N^{1/2}$), for general functions a weaker attack $N^3 < O(S^3 \cdot T)$ (e.g. $S = T \approx N^{3/4}$) exists.

The best lower bounds are of the form $N > \tilde{\Omega}(S \cdot T)$ and hold for random permutations and functions.

Motivated by the application to proofs of space (PoSpace), we construct functions for which we can prove much better lower bounds of the form $N^k > \tilde{\Omega}(S^k \cdot T)$ (for any constant k). Our construction does not contradict the existing attacks, as these attacks require that the function to be inverted can be efficiently computed in forward direction. For the application to PoSpace it is sufficient that the entire function table can be computed in time quasilinear in N.

The simplest function that beats the existing bound is build from a random function $g : [N] \times [N] \to [N]$ and a random permutations $f, f' : [N] \to [N]$ and is defined as $h(x) = g(x, x')$ where $f(x) = f(x') + 1$ (instead of $+1$ one can use any other bijection without fixpoints). For this function we prove a lower bound of $N^2 < O(S^2 \cdot T)$. Note that h cannot be efficiently evaluated on input x as one has on find $x' = f^{-1}(f(x) - 1)$, but its function table can be computed in time $O(N)$ by first computing the function table for f^{-1}.

3.13 Integer Commitments

David Pointcheval (ENS – Paris, FR)

Committing integers and proving relations between them is an essential ingredient in many cryptographic protocols. Among them, range proofs have shown to be fundamental. They consist in proving that a committed integer lies in a public interval. By the way, it can also be seen as a particular case of the more general Diophantine relations: for the committed vector of integers $\vec{x}$, there exists a vector of integers $\vec{w}$ such that $P(\vec{x}, \vec{w}) = 0$, where P is a polynomial.

In this talk, we revisit the security strength of the statistically hiding commitment scheme over the integers due to Damgård-Fujisaki, and the zero-knowledge proofs of knowledge of openings.

First, we show how to remove the Strong RSA assumption and replace it by the standard RSA assumption in the security proofs. This improvement naturally extends to generalized commitments and more complex proofs without modifying the original protocols.

Second, we design an interactive technique turning commitment scheme over the integers into commitment scheme modulo a prime p. Still under the RSA assumption, this results in more efficient proofs of relations between the committed values. Our methods thus improve upon existing proof systems regarding Diophantine relations both in terms of performance and security.

We illustrate that with more efficient range proofs under the sole RSA assumption.

3.14 Securing Public Key Encryption in the Presence of Bad Randomness

Jacob Schuldt (AIST – Tsukuba, JP)

In this talk, we firstly motivate the need for encryption secure in the presence of bad randomness, and revisit the notion of related randomness security by Paterson, Schuldt, and Sibborn, as well as some of the known constructions of related randomness secure encryption. We then highlight an inherent limitation of the related randomness security notion: if the family of related randomness functions is sufficiently rich to express the encryption function of the considered scheme, then security cannot be achieved. This might help explain why the previous standard model constructions only achieve security for polynomial function families.

To address this limitation, we propose a new notion, related refreshable randomness security, which captures that an adversary has limited time to attack a system before new entropy is added. In this setting, we construct an encryption scheme which remains secure in the standard model for arbitrary function families of size 2^p (where p is polynomial in the security parameter) that satisfy certain collision-resistant and output-unpredictability properties. This captures a rich class of functions, which includes, as a special case, circuits of polynomial size.

3.15 On the Impossibility of Tight Cryptographic Reductions

Sven Schäge (Ruhr-Universität Bochum, DE)

The existence of tight reductions in cryptographic security proofs is an important question, motivated by the theoretical search for cryptosystems whose security guarantees are truly independent of adversarial behavior and the practical necessity of concrete security bounds for the theoretically-sound selection of cryptographic parameters. At Eurocrypt 2002, Coron described a meta-reduction technique that allows to prove the impossibility of tight reductions for certain digital signature schemes. This seminal result has found many further interesting applications. However, due to a technical subtlety in the argument, the applicability of this

technique beyond digital signatures in the single-user setting has turned out to be rather limited.

We describe a new meta-reduction technique for proving such impossibility results, which improves on known ones in several ways. First, it enables interesting novel applications. This includes a formal proof that for certain cryptographic primitives (including public-key encryption/key encapsulation mechanisms and digital signatures), the security loss incurred when the primitive is transferred from an idealized single-user setting to the more realistic multi-user setting is impossible to avoid, and a lower tightness bound for non-interactive key exchange protocols. Second, the technique allows to rule out tight reductions from a very general class of non-interactive complexity assumptions. Third, the provided bounds are quantitatively and qualitatively better, yet simpler, than the bounds derived from Coron's technique and its extensions.

3.16   Android Security using Static Analysis Techniques

Suzanna Schmeelk (Columbia University – New York, US)

Static analysis is a traditional technique for software transformation and analysis. It has also become a means to detect cyber security vulnerabilities and malware and recently has been extended to the mobile-computing arena for security-related analyses. This talk examines over fifty recent security papers that are published in top conferences, journals and technical reports and characterizes the current research. The papers were selected based on either their high citings by other top research or they introduced either a novel analysis technique or a novel security issue analysis. Our research systematically constructs a static analysis landscape by charting and characterizing analysis strengths and limitations in both accuracy and security threats. It identifies two types of static analysis motivations which affect the soundness of an analysis methodology: (1) techniques for analyzing software for vulnerabilities and (2) techniques used to examine applications for malware, which may lead to malware mitigation. We analyze techniques and tools for effort-level required use by security analysists and connect the reported static analysis motivations to both Mitre's attack taxonomy as well as Mitre's vulnerability taxonomy to aid completeness. Our findings include identifying vulnerabilities which are not being systematically researched, identifying best practices for developers and characterizing technique usability metrics for integrating the analysis into a security analysis process.

3.17 The OPTLS Protocol and TLS 1.3

Hoeteck Wee (ENS – Paris, FR)

We present the OPTLS key-exchange protocol, its design, rationale and cryptographic analysis. OPTLS design has been motivated by the ongoing work in the TLS working group of the IETF for specifying TLS 1.3, the next-generation TLS protocol. The latter effort is intended to revamp the security of TLS that has been shown inadequate in many instances as well as to add new security and functional features. The main additions that influence the cryptographic design of TLS 1.3 (hence also of OPTLS) are a new "0-RTT requirement" (0-RTT stands for "zero round trip time") to allow clients that have a previously retrieved or cached public key of the server to send protected data already in the first flow of the protocol; making forward secrecy (PFS) a mandatory requirement; and moving to elliptic curves as the main cryptographic basis for the protocol (for performance and security reasons). Accommodating these requirements calls for moving away from the traditional RSA-centric design of TLS in favor of a protocol based on Diffie-Hellman techniques. OPTLS offers a simple design framework that supports all the above requirements with a uniform and modular logic that helps in the specification, analysis, performance optimization, and future maintenance of the protocol. An earlier (draft) specification of TLS 1.3 built upon the OPTLS framework as a basis for the cryptographic core of the handshake protocol, adapting the different modes of OPTLS and its HKDF-based key derivation to the TLS 1.3 context.

Participants

- Adekunle Oluseyi Afolabi
University of Kuopio, FI
- Fabrice Benhamouda
IBM Thomas J. Watson Research
Center – Yorktown Heights, US
- Johannes A. Buchmann
TU Darmstadt, DE
- David Cash
Rutgers University, US
- Pooya Farshim
ENS – Paris, FR
- Marc Fischlin
TU Darmstadt, DE
- Pierre-Alain Fouque
University of Rennes, FR
- Vipul Goyal
Microsoft Research India –
Bangalore, IN
- Iftach Haitner
Tel Aviv University, IL
- Dennis Hofheinz
KIT – Karlsruher Institut für
Technologie, DE

- Antoine Joux
CNRS and University Pierre &
Marie Curie – Paris, FR
- Eike Kiltz
Ruhr-Universität Bochum, DE
- Alexander Koch
KIT – Karlsruher Institut für
Technologie, DE
- Tal Malkin
Columbia Univ. – New York, US
- Alexander May
Ruhr-Universität Bochum, DE
- Jörn Müller-Quade
KIT – Karlsruher Institut für
Technologie, DE
- Phong Q. Nguyen
Inria and CNRS/JFLI, FR, and
University of Tokyo, JP
- Kenneth G. Paterson
Royal Holloway University of
London, GB

- Krzysztof Pietrzak
IST Austria –
Klosterneuburg, AT
- David Pointcheval
ENS – Paris, FR
- Tal Rabin
IBM Thomas J. Watson Research
Center – Yorktown Heights, US
- Sven Schäge
Ruhr-Universität Bochum, DE
- Suzanna Schmeelk
Columbia Univ. – New York, US
- Dominique Schröder
Univ. Erlangen-Nürnberg, DE
- Jacob Schuldt
AIST – Tsukuba, JP
- Vinod Vaikuntanathan
MIT – Cambridge, US
- Hoeteck Wee
ENS – Paris, FR