

MATHEMATISCHES FORSCHUNGSINSTITUT OBERWOLFACH

Report No. 35/2014

DOI: 10.4171/OWR/2014/35

Cryptography

Organised by
Johannes Buchmann, Darmstadt
Shafi Goldwasser, Cambridge MA/Rehovot

27 July – 2 August 2014

ABSTRACT. The Oberwolfach workshop *Cryptography* brought together scientists from cryptography with mathematicians specializing in the algorithmic problems underlying cryptographic security. The goal of the workshop was to stimulate interaction and collaboration that enables a holistic approach to designing cryptography from the mathematical foundations to practical applications. The workshop covered basic computational problems such as factoring and computing discrete logarithms and short vectors. It addressed fundamental research results leading to innovative cryptography for protecting security and privacy in cloud applications. It also covered some practical applications.

Mathematics Subject Classification (2010): 68M11, 68Q25, 68W40.

Introduction by the Organisers

The goal of the workshop *Cryptography*, organized by Johannes Buchmann (Darmstadt) and Shafi Goldwasser (Boston) was to stimulate interaction and collaboration between mathematicians and computer scientists that enables a holistic approach to designing cryptography from the mathematical foundations to practical applications.

The goal of the workshop is of great importance both from a research and an application point of view. Cryptographic schemes are of very important in practice as they are indispensable building blocks of cyber security solutions. On the other hand, the development of cryptography is a major scientific challenge since its security is threatened by new attacks, for example by quantum computers, and the protection of new applications such as cloud computing requires innovative cryptographic models and techniques.

The talks given at the workshop covered important recent results in the areas relevant for the workshop.

The talks on the mathematical foundations addressed both traditional and more recent algorithmic problems that serve as the security basis of modern cryptography. Among the topics were new algorithms for factoring integers and computing discrete logarithms in the multiplicative group of finite fields, problems whose hardness is the foundation of current public-key cryptography. The presented results showed that there is progress in dealing with these problems. This implies that alternative problems must be studied. Progress with respect to a very important alternative, the problem of finding short vectors in lattices (SVP), was the topic of several talks at the workshop. SVP not only allows for the construction of new encryption and signature schemes but is also the basis of most relevant advanced cryptographic constructions. The presentations at the conference illuminated the mathematical structure of SVP and addressed new algorithmic approaches.

The presentations on advanced cryptographic constructions covered both security and functionality aspects. For example, one of the security topics was “quantum random oracles”. In view of the development of quantum computers it is a fundamental task of cryptography to take such computers into account when modelling security. The talks on cryptography with advanced functionality covered homomorphic encryption and signatures, program obfuscation, and garbled RAMs. These techniques address the very important problem of privacy protection when data storage and computation on these data is outsourced. The recent breakthroughs in this area were very well represented in the talks of the workshop.

As cryptography is a topic with great real world relevance, there were finally a number of talks that dealt with practical issues such as differential privacy, multi-party-computation, and the practical exploitability of the TLS protocol.

The talks stimulated a very intense discussion among the participating scientists from the different fields. This discussion lead already to further collaborations. At the end of the workshop, the participants were very enthusiastic about its success. They expressed the hope that the workshop will be repeated in the not too far future.

Acknowledgement: The MFO and the workshop organizers would like to thank the National Science Foundation for supporting the participation of junior researchers in the workshop by the grant DMS-1049268, “US Junior Oberwolfach Fellows”.

Workshop: Cryptography**Table of Contents**

Antoine Joux	
<i>A simplified setting for discrete logarithms in small characteristic finite fields</i>	1939
Benny Applebaum (joint with Jonathan Avron, Christina Brzuska)	
<i>Arithmetic Cryptography</i>	1940
Yuval Ishai (joint with Daniel Genkin, Manoj Prabhakaran, Amit Sahai, Eran Tromer)	
<i>Circuits Resilient to Additive Attacks with Applications to Secure Computation</i>	1941
Elette Boyle (joint with Kai-Min Chung, Rafael Pass)	
<i>Large-Scale Secure Computation</i>	1942
Daniele Micciancio (joint with Michael Walter)	
<i>A New Variant of Kannan's Lattice Enumeration Algorithm</i>	1944
Sergey Gorbunov (joint with Vinod Vaikuntanathan, Daniel Wichs)	
<i>(Leveled) Fully Homomorphic Signatures from Lattices</i>	1944
Jintai Ding (joint with Chengdong Tao)	
<i>A New Algorithm for Solving the Approximate Common Divisor Problem and Cryptanalysis of the Fully Homomorphic Encryption Schemes</i>	1946
Dennis Hofheinz	
<i>Compact and tightly secure cryptography in the standard model</i>	1948
Chris Peikert	
<i>Ring switching and Bootstrapping Fully Homomorphic Encryption</i>	1950
Moni Naor (joint with Ben Fisch, Daniel Freund)	
<i>Physical zero knowledge</i>	1950
Özgür Dagdelen (joint with Sebastian Gajek, Florian Göpfert)	
<i>Learning with Errors in the Exponent</i>	1952
Iftach Haitner (joint with Eliad Tzfatia)	
<i>An Almost-Optimally Fair Three-Party Coin-Flipping Protocol</i>	1953
Huijia Lin (joint with Irit Dinur, Shafi Goldwasser)	
<i>The Computational Benefit of Correlated Instances</i>	1955
Alon Rosen (joint with Vipul Goyal, Silas Richelson, Margarita Vald)	
<i>An Algebraic Approach to Non-Malleability</i>	1956

Marc Fischlin (joint with Özgür Dagdelen, Tommaso Gagliardoni)	
<i>The Fiat-Shamir Transformation in the Quantum Random Oracle Model</i>	1957
Steven Galbraith	
<i>Lattice Algorithms for Learning with Errors</i>	1959
Guy Rothblum (joint with Salil Vadhan, Avi Wigderson)	
<i>Interactive Proofs of Proximity: Delegating Computation in Sublinear Time</i>	1961
Yael Tauman Kalai, Ron D. Rothblum	
<i>How to Delegate Computations: The Power of No-Signaling Proofs</i>	1961
Rafail Ostrovsky (joint with Sanjam Garg, Steve Lu, Alessandra Scafuro)	
<i>Garbled RAM from One-way Functions</i>	1963
Hendrik W. Lenstra, Alice Silverberg	
<i>Lattices with symmetry, and the extended tensor algebra</i>	1966
Stefano Tessaro (joint with David Cash)	
<i>The Locality of Searchable Symmetric Encryption</i>	1967
Cynthia Dwork (joint with Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, Aaron Roth)	
<i>A Surprising Application of Differential Privacy</i>	1968
Omer Paneth (joint with Nir Bitansky, Ran Canetti, Yael Tauman Kalai)	
<i>On Virtual Grey Box Obfuscation for General Circuits</i>	1970
Rafael Pass (joint with Karn Seth, Sidharth Telang)	
<i>Indistinguishability Obfuscation from Semantically-Secure Multilinear Encodings</i>	1972
Mark Zhandry	
<i>How to Avoid Obfuscation Using Witness PRFs</i>	1974
Daniel Wichs (joint with Craig Gentry, Shai Halevi, Mariana Raykova)	
<i>Outsourcing Private RAM Computation</i>	1977
Nigel P. Smart	
<i>Practical Multi-Party Computation</i>	1978
Claus P. Schnorr	
<i>Factoring Integers by CVP Algorithms</i>	1981
Vipul Goyal	
<i>Non-Black-Box Simulation in the Fully Concurrent Setting</i>	1982
Tanja Lange (joint with S. Checkowoy, R. Nicolierhagen, M. Fredrikson, A. Everspaugh, M. Green, T. Ristenport, J. Moskiwicz, and H. Shochom)	
<i>On the practical exploitability of Dual EC DRBG in TLS implementations</i>	1984

Sergey Gorbunov (joint with Dan Boneh, Craig Gentry, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, Dhinakaran Vinayagamurthy) <i>Fully Key-Homomorphic Encryption, Arithmetic Circuit ABE, and Compact Garbled Circuits</i>	1985
--	------

Abstracts

A simplified setting for discrete logarithms in small characteristic finite fields

ANTOINE JOUX

The hardness of computing discrete logarithms in finite field has served as a foundation for many public key cryptosystems. In the last two years, tremendous progress have been made in the case of small characteristic finite fields.

In this talk, we present a simplified description of the algorithmic framework that has been developed to solve this problem faster. This framework is an index calculus approach that relies on two main ingredients, the definition of the extension field and the generation of multiplicative relations in this field. Given a base field \mathbb{F}_q , we construct its extension field \mathbb{F}_{q^k} in the following way: we find two polynomials of low degree h_0 and h_1 with coefficients in F_q such that $x^q h_1(x) - h_0(x)$ has an irreducible factor I_k of degree k over \mathbb{F}_q . Let θ denotes a root of I_k and define \mathbb{F}_{q^k} as $\mathbb{F}_q(\theta)$. In the larger finite field, we know that by construction θ satisfy the relation:

$$(1) \quad \theta^q = \frac{h_0(\theta)}{h_1(\theta)}.$$

Note that, it is also possible to work with an alternative definition of the form $\theta = h_0(\theta^q)/h_1(\theta^q)$.

To generate relations, we start from the well-known identity:

$$(2) \quad x^q - x = \prod_{\alpha \in \mathbb{F}_q} x - \alpha.$$

Replacing x by $A(\theta)/B(\theta)$ in (2) and multiplying by $B(\theta)^q$ we find:

$$B(\theta)A(\theta)^q - A(\theta)B(\theta)^q = B(\theta) \prod_{\alpha \in \mathbb{F}_q} (A(\theta) - \alpha B(\theta)).$$

Assume that A and B are two polynomials of degree at most D with coefficients in \mathbb{F}_q , by linearity of the Frobenius, we can replace $A(\theta)^q$ by $A(h_0(\theta)/h_1(\theta))$ and rewrite the equation as:

$$(3) \quad \frac{[A, B]_D(\theta)}{h_1(\theta)^D} = \prod_{\alpha \in \mathbb{P}_1(\mathbb{F}_q)} (A(\theta) - \alpha B(\theta)).$$

To make (3) more compact, we define $A(\theta) - \alpha B(\theta)$ as an alias for $B(\theta)$ when α is the point at infinity in $\mathbb{P}_1(\mathbb{F}_q)$. Moreover, we let $[A, B]_D$ denote the polynomial $h_1(x)^D (B(x)A(h_0(x)/h_1(x)) - A(x)B(h_0(x)/h_1(x)))$. We remark that $[A, B]_D$ is a polynomial of degree at most $D \cdot (H + 1)$ where H denotes the maximum of the degrees of h_0 and h_1 . The bracket $[A, B]_D$ has some interesting properties: it is alternating and \mathbb{F}_q bilinear. Using these properties, we can easily transform A and B to make them monic and ensure that $\deg(A) > \deg(B)$. In particular, this

indicates that choosing $A(X) = X^D + A_0(X)$ and $B(X) = X^{D-1} + B_0(X)$ with A_0 and B_0 of degree $D - 2$ is a good way to obtain distinct equations.

Finally, remark that when $[A, B]_D$ factors into terms of degree at most D , then (3) gives a multiplicative relation between monic irreducible polynomials of degree at most D . Since there are approximately q^D/D such polynomials and since the above process generates at most $q^{2(D-1)}$ equations, it is clear that $D = 0$ and $D = 1$ cannot be enough. Taking $D = 3$, we obtain an algorithm with complexity $O(q^7)$ to obtain the logarithms of the low degree irreducibles. This is in line with the literature on the topic.

Once this is done, we use a descent procedure to recursively express any element of the finite field \mathbb{F}_{q^k} into elements represented by polynomials of lower degree. This procedure is quite complex but ultimately leads to a quasi-polynomial time algorithm for the discrete logarithm problem in small characteristic finite fields.

REFERENCES

- [BGJT14] Razvan Barbulescu, Pierrick Gaudry, Antoine Joux, and Emmanuel Thomé. A heuristic quasi-polynomial algorithm for discrete logarithm in finite fields of small characteristic. In *EUROCRYPT*, pages 1–16, 2014.
- [GGMZ13a] Faruk Göloğlu, Robert Granger, Gary McGuire, and Jens Zumbärgel. On the function field sieve and the impact of higher splitting probabilities - application to discrete logarithms in $\mathbb{F}_{2^{1971}}$ and $\mathbb{F}_{2^{3164}}$. In *CRYPTO (2)*, pages 109–128, 2013.
- [Jou13a] Antoine Joux. Faster index calculus for the medium prime case application to 1175-bit and 1425-bit finite fields. In *EUROCRYPT*, pages 177–193, 2013.
- [Jou14a] Antoine Joux. A new index calculus algorithm with complexity $L(1/4+o(1))$ in very small characteristic. In *Selected Areas in Cryptography-SAC 2013*, volume 8282 of *Lecture Notes in Computer Science*, pages 355–382. Springer, 2014.

Arithmetic Cryptography

BENNY APPLEBAUM

(joint work with Jonathan Avron, Christina Brzuska)

We study the possibility of computing cryptographic primitives in a fully-black-box arithmetic model over a finite field F . In this model, the input to a cryptographic primitive (e.g., encryption scheme) is given as a sequence of field elements, the honest parties are implemented by arithmetic circuits which make only a black-box use of the underlying field, and the adversary has a full (non-black-box) access to the field. This model captures many standard information-theoretic constructions including the classical secure multiparty protocols of [BGW88, CCD88].

We prove several positive and negative results in this model for various cryptographic tasks. On the positive side, we show that, under coding-related assumptions, computational primitives like commitment schemes, public-key encryption, oblivious transfer, and general secure two-party computation can be implemented in the arithmetic model. On the negative side, we prove that garbled circuits, multiplicative homomorphic encryption, and secure computation with low online communication complexity cannot be achieved in this model.

Our results reveal a qualitative difference between the standard model and the arithmetic model, and explain, in retrospect, some of the limitations of previous constructions.

REFERENCES

- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation (extended abstract). pages 1–10, 1988.
- [CCD88] David Chaum, Claude Crépeau, and Ivan Damgård. Multiparty unconditionally secure protocols (extended abstract). pages 11–19, 1988.

Circuits Resilient to Additive Attacks with Applications to Secure Computation

YUVAL ISHAI

(joint work with Daniel Genkin, Manoj Prabhakaran, Amit Sahai, Eran Tromer)

We study the question of protecting arithmetic circuits against *additive* attacks, which can add an arbitrary fixed value to each wire in the circuit. This extends the notion of algebraic manipulation detection (AMD) codes, which protect *information* against additive attacks, to that of *AMD circuits* which protect *computation*.

We present a construction of such AMD circuits: any arithmetic circuit C over a finite field \mathbb{F} can be converted into a functionally-equivalent randomized arithmetic circuit \widehat{C} of size $O(|C|)$ that is fault-tolerant in the following sense. For any additive attack on the wires of \widehat{C} , its effect on the output of \widehat{C} can be simulated, up to $O(|C|/|\mathbb{F}|)$ statistical distance, by an additive attack on just the input and output. Given a small tamper-proof encoder/decoder for AMD codes, the input and output can be protected as well.

We also give an alternative construction, applicable to small fields (for example, to protect Boolean circuits against wire-toggling attacks). It uses a small tamper-proof decoder to ensure that, except with negligible failure probability, either the output is correct or tampering is detected.

Our study of AMD circuits is motivated by the goal of simplifying and improving protocols for secure multiparty computation (MPC). Typically, securing MPC protocols against *active* adversaries is much more difficult than securing them against *passive* adversaries, who follow the protocol but try to learn additional information from messages they receive. We observe that in simple MPC protocols that were designed to protect circuit evaluation only against *passive* adversaries, the effect of any *active* adversary corresponds precisely to an additive attack on the original circuit's wires. Thus, to securely evaluate a circuit C in the presence of active adversaries, it suffices to apply the passive-secure protocol to \widehat{C} . We use this methodology to simplify feasibility results and attain efficiency improvements in several standard MPC models.

Our work gives rise to several open questions and directions for further research. The first is to improve the security and efficiency of the construction for

the case of small fields. In our current construction, there may be a correlation between the input and the event that an error is detected. Such correlations prevent us from applying the constructions over small fields in the context of MPC. A second direction is to extend the study of AMD circuits to accommodate other classes of attacks. Finally, the applications to MPC can be extended to cover information-theoretic “constant-rate” protocols with sub-optimal security threshold and computationally secure protocols.

REFERENCES

- [1] D. Genkin, Y. Ishai, M. Prabhakaran, A. Sahai, E. Tromer, *Circuits Resilient to Additive Attacks with Applications to Secure Computation*, Proceedings of STOC 2014, 495–504.

Large-Scale Secure Computation

ELETTE BOYLE

(joint work with Kai-Min Chung, Rafael Pass)

The notion of secure multi-party computation (MPC), introduced in the seminal works of Yao and Goldreich, Micali and Wigderson, is one of the cornerstones in cryptography. An MPC protocol for computing a function f allows a group of parties to jointly evaluate f over their private inputs, with the property that an adversary who corrupts a subset of the parties does not learn anything beyond the inputs of the corrupted parties and the output of the function f .

An emerging area of potential applications for secure MPC is to address privacy concerns in data aggregation and analysis to match the explosive current growth of the amount of available data. Cryptographic techniques such as MPC for secure function evaluation where *data items are equated with servers* can be utilized to prevent unnecessary leakage of information.

However, before MPC can be effectively used to address today’s challenges, we need protocols whose efficiency and communication requirements scale practically to the modern regime of massive data. When the data set contains tens of thousands of users’ web traffic patterns or personal genetic information, it becomes unreasonable to assume any single user can provide memory, computation, or communication resources on the order of the data of *all users*. When the computations to be executed are lightweight, depend on a small subset of inputs, or require small memory, it will be unacceptable to obliterate these savings to achieve security. In this regime, the efficiency of existing solutions breaks down: either requiring resources linear in the *circuit representation* size of the function (including works in the line of scalable MPC [1, 2, 3, 4, 5]), or requiring parties to store and communicate information on the order of *all parties’* combined inputs (by naïve extension of two-party protocols for RAM programs such as [6, 7, 8, 9] to the multiparty setting).

We achieve secure multiparty computation that directly supports evaluating RAM programs on parties’ inputs, with a protocol that preserves the *per-party* memory and computation complexity requirements of the participating parties.

Our construction is information theoretically secure against $(1/3 - \epsilon)$ statically scheduled corruptions, within a synchronous communication network.

Theorem (*Informal – Load-Balanced, Communication-Local MPC for RAM Programs*) For any constant $\epsilon > 0$, there exists an n -party statistically secure (with error negligible in n) protocol for computing any adaptively chosen sequence of N RAM programs Π_j (that may have shared state), handling $(1/3 - \epsilon)$ fraction static corruptions making an initial use of a single broadcast per party (of $\text{polylog}(n)$ bits), and with the following complexities (where $|x|, |y|$ denote input and output size):

- Memory per party: $\tilde{O}(|x| + \max_{j=1}^N \text{Space}(\Pi_j)/n)$.
- Computation per party: $\tilde{O}(|x| + \sum_{j=1}^N \text{Time}(\Pi_j)/n + N|y|)$.
- Round complexity: $\tilde{O}(\sum_{j=1}^N \text{Time}(\Pi_j))$.

where $\text{Space}(\Pi)$ and $\text{Time}(\Pi)$ denote the worst-case space and runtime requirements of Π over different inputs. Additionally, our protocol achieves $\text{polylog}(n)$ communication locality, and a strong “online” load-balancing guarantee such that at *all times* during the protocol, all parties’ communication and computation loads vary by at most a multiplicative factor of $\text{polylog}(n)$ (up to a $\text{polylog}(n)$ additive term).

Note that the initial one-time uses of a broadcast channel can be implemented via execution of a broadcast protocol of choice. We separate the broadcast cost from our protocol complexity measures to emphasize that any (existing or future) broadcast protocol can be directly plugged in, yielding associated desirable properties.

We additionally remark that even without considering communication locality, amortization of setup over multiple RAM program executions, or the communication/computation load-balancing properties achieved by our protocol, our results already yield the first protocol whose *total* communication and computation complexities for securely evaluating a *single* RAM program Π grow as $\text{poly}(n) + \tilde{O}(\text{Time}(\Pi))$ while simultaneously requiring only $\tilde{O}(|x| + \text{Space}(\Pi)/n)$ memory per party. Indeed, all existing solutions either require converting the program into a circuit representation, or require parties to maintain storage $\Omega(n|x|)$.

REFERENCES

- [1] Ivan Damgård and Yuval Ishai. Scalable secure multiparty computation. In *CRYPTO*, pages 501–520, 2006.
- [2] Ivan Damgård and Jesper Buus Nielsen. Scalable and unconditionally secure multiparty computation. In *CRYPTO*, pages 572–590, 2007.
- [3] Ivan Damgård, Yuval Ishai, Mikkel Krøigaard, Jesper Buus Nielsen, and Adam Smith. Scalable multiparty computation with nearly optimal work and resilience. In *CRYPTO*, pages 241–261, 2008.
- [4] Ivan Damgård, Yuval Ishai, and Mikkel Krøigaard. Perfectly secure multiparty computation and the computational overhead of cryptography. In *EUROCRYPT*, pages 445–465, 2010.
- [5] Varsha Dani, Valerie King, Mahnush Movahedi, and Jared Saia. Breaking the $o(nm)$ bit barrier: Secure multiparty computation with a static adversary. *CoRR*, abs/1203.0289, 2012.

- [6] Rafail Ostrovsky and Victor Shoup. Private information storage (extended abstract). In *STOC*, pages 294–303, 1997.
- [7] S. Dov Gordon, Jonathan Katz, Vladimir Kolesnikov, Tal Malkin, Mariana Raykova, and Yevgeniy Vahlis. Secure computation with sublinear amortized work. In *ACM Conference on Computer and Communications Security*, pages 513–524, 2012.
- [8] Steve Lu and Rafail Ostrovsky. Distributed oblivious RAM for secure two-party computation. In *TCC*, pages 377–396, 2013.
- [9] Craig Gentry, Kenny A. Goldman, Shai Halevi, Charanjit S. Jutla, Mariana Raykova, and Daniel Wichs. Optimizing ORAM and using it efficiently for secure computation. In *Privacy Enhancing Technologies*, pages 1–18, 2013.

A New Variant of Kannan’s Lattice Enumeration Algorithm

DANIELE MICCIANCIO

(joint work with Michael Walter)

Enumeration algorithms are the best currently known methods to solve lattice problems, both in theory (within the class of polynomial space algorithms) and in practice (where they are routinely used to evaluate the concrete security of lattice cryptography). However, there is a big gap between our theoretical understanding and the practical performance of lattice enumeration algorithms. We present a variant of the algorithm of Kannan, matching its theoretical asymptotic performance, but with much smaller overhead, comparable to the algorithms used in practice already in relatively small dimension.

(Leveled) Fully Homomorphic Signatures from Lattices

SERGEY GORBUNOV

(joint work with Vinod Vaikuntanathan, Daniel Wichs)

With advances in cloud computing, an increasing amount of sensitive data is stored and computations on them are performed remotely, raising questions of privacy of the data and correctness of computations. Recently, a number of cryptographic schemes have been developed to address these concerns. For example, fully homomorphic encryption [1, 2] enables us to compute on encrypted data, paving the road to achieving privacy in outsourcing. Many flavors of verifiable outsourcing schemes have been developed to deal with the question of correctness of computations (cf. [3, 4, 5, 6] and many others). A particularly natural way to verifiably outsource computation is through the notion of homomorphic signatures [7, 8, 9].

A homomorphic signature scheme is one where anyone can homomorphically compute on the signatures $\vec{\sigma} = (\sigma_1, \sigma_2, \dots, \sigma_\ell)$ corresponding to a dataset $\vec{\mu} = (\mu_1, \mu_2, \dots, \mu_\ell)$ and produce a signature σ' for a circuit C and the result $\mu' = C(\vec{\mu})$ of applying C to the dataset $\vec{\mu}$. Given only the public key pk and the signature σ' on the circuit C and a message μ' , anyone can verify that σ' is indeed the result of applying C to some set of signed messages $\vec{\mu}$. In order to tie the signature to a particular dataset, we “tag” each dataset of messages, and give the tag to the

verification algorithm as well. A key feature is that this verification can be done without knowing the original dataset $\vec{\mu}$.

The signature σ' “proves” that the computation was done correctly, in the sense that computing a signature σ' for any pair (C, μ') where $\mu' \neq C(\vec{\mu})$ is hard for any PPT adversary. Moreover, the resulting signature is compact, namely, its size and the time to verify it depends neither on the size of the original data or the size of the circuit that was computed on it. This gives us a very natural, publicly verifiable scheme to outsource computation (in an amortized setting).

However, constructions of homomorphic signatures have been few and far between. In particular:

- The initial schemes [7, 8] handled only linear functions. The state of the art is a scheme of Boneh and Freeman [9] that can compute constant degree polynomial functions on signed messages.
- The schemes are shown secure in the random oracle model.
- Finally, the polynomially homomorphic schemes rely on the short integer solutions (SIS) problem on *ideal lattices*. In contrast, in the case of fully homomorphic *encryption*, we know several solutions by now that rely on the SIS problem on *arbitrary lattices* with no ideal structure.

In this work, we construct the first *leveled fully homomorphic signature* schemes that can evaluate arbitrary circuits over signed data, where only the maximal depth d of the circuit needs to be fixed a priori. The size of the evaluated signature grows polynomially in d , but is otherwise independent of the circuit size or the data size. Our solutions are based on the hardness of the *small integer solution* (SIS) problem, which is in turn implied by the worst-case hardness of problems in standard lattices. We get a scheme in the standard model, albeit with large public parameters whose size must exceed the total size of all signed data. In the random-oracle model, we get a scheme with short public parameters.

As a building block of independent interest, we introduce a new notion called *homomorphic trapdoor functions* (HTDF). We show to how construct homomorphic signatures using HTDFs as a black box. We construct HTDFs based on the SIS problem by relying on a recent technique developed by Boneh et al. [10] in the context of attribute-based encryption.

Interesting open problems include removing the dependency on the circuit depth in the evaluated signature and making the size of the public parameters independent on the size of the dataset in the standard model.

REFERENCES

- [1] Craig Gentry. Fully homomorphic encryption using ideal lattices. In STOC, pages 169-178, 2009.
- [2] Zvika Brakerski, Craig Gentry, and Vinod Vaikuntanathan. (leveled) fully homomorphic encryption without bootstrapping. In ITCS, pages 309-325, 2012.
- [3] Silvio Micali. Computationally sound proofs. SIAM J. Comput., 30(4):1253-1298, 2000.
- [4] Joe Kilian. A note on efficient zero-knowledge proofs and arguments (extended abstract). In STOC, 1992.

- [5] Shafi Goldwasser, Yael Tauman Kalai, and Guy N. Rothblum. Delegating computation: interactive proofs for muggles. In *STOC*, pages 113-122, 2008.
- [6] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. Delegation for bounded space. In *STOC*, 2013.
- [7] Denis Xavier Charles, Kamal Jain, and Kristin Lauter. Signatures for network coding. *IJI-CoT*, 2009.
- [8] Dan Boneh, David Mandell Freeman, Jonathan Katz, and Brent Waters. Signing a linear subspace: Signature schemes for network coding. In *Public Key Cryptography*, 2009.
- [9] Dan Boneh and David Mandell Freeman. Homomorphic signatures for polynomial functions. In *EUROCRYPT*, 2011.
- [10] Dan Boneh, Craig Gentry, Sergey Gorbunov, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, and Dhinakaran Vinayagamurthy. Fully key-homomorphic encryption, arithmetic circuit abe and compact garbled circuits. In *EUROCRYPT*, 2014.

A New Algorithm for Solving the Approximate Common Divisor Problem and Cryptanalysis of the Fully Homomorphic Encryption Schemes

JINTAI DING

(joint work with Chengdong Tao)

Approximate common divisor problem was first introduced by Howgrave-Graham in [14]. There are several fully homomorphic encryption schemes which are based on the approximate common divisors problem [8][9][10][19]. Approximate common divisors problem is defined as follows:

- **General approximate common divisors(GACD) problem:** *For a set of parameters γ , η , and ρ , given polynomial (in γ , η , and ρ) many different integers in the form: $x_i = pq_i + r_i (i = 1, \dots, n)$, the problem is to recover p , where x_i are of bit length γ , p is of bit length η , r_i are small integers with the bit length no more than ρ . Here r_i are called the error terms.*

A simple approach for solving GACD problem is exhaustive search on the error terms. If r_i are sufficiently small, then we can find p by exhaustive search, i.e., one can try every r_1 and r_2 and check whether $\gcd(x_1 - r_1, x_2 - r_2)$ is sufficiently large and eventually recover p , where $\gcd()$ is the algorithm for solving the greatest common divisor. The state of the art algorithm for computing greatest common divisor is the Stehlè-Zimmermann algorithm with time complexity $O(\gamma)$ for integers of γ bits[18]. In *EUROCRYPT'12*, Chen and Nguyen gave an algorithm which provides an exponential speedup over exhaustive search to solve approximate common divisors problem [6], which is essentially based a clever exhaustive search on the error terms through certain polynomials. However, their approach requires very large memories. For their algorithm, they only need 2 elements in the set of x_i and the complexity is given as $\mathcal{O}(2^{\frac{3}{2}\rho\gamma})$.

In [14], Howgrave-Graham also gives a lattice approach to solve two elements GACD problem. This approach is related to Coppersmith's algorithm for finding small solutions to univariate and bivariate modular equations. When $\frac{\rho}{\gamma}$ is smaller

than $(\frac{\eta}{\gamma})^2$, this approach recovers p . However, when ρ, η, γ do not satisfy the constraint, the approach does not degrade gracefully. Furthermore, in [7], Cohn and Heninger analyze the multivariate generalization of Howgrave-Graham's algorithm for the GACD problem by using many x_i . In this algorithm, the GACD problem used in cryptography is reduced to running the LLL algorithm on a lattice basis of high dimension and large entries to directly find all the error terms r_i . However, in [6], they show that the Cohn-Heninger attack on the FHE challenges in [8] is actually slower than exhaustive search on the challenges, and therefore much slower than the attack in [6].

In this talk, we propose a new algorithm for solving the approximate common divisors problem. We consider the first t integers x_1, \dots, x_t in GACD problem, where t satisfies $\frac{\gamma}{t} + \rho + t \log \delta + \frac{3}{2} \log t + 2 < \eta$, and δ is the root Hermit factor. Without loss of generality, we assume that $x_t = \max\{x_1, \dots, x_t\}$. Since the error terms r_1, \dots, r_t are small relative to p , the main ideal of our algorithm is to find a vector $\mathbf{u} = (u_1, \dots, u_t)$ such that $\sum_{i=1}^t u_i \cdot x_i < p$ and $\|\mathbf{u}\| < 2^{\eta-\rho-1}/\sqrt{t}$. Assume that we find such a vector \mathbf{u} , then we obtain an equation over \mathbb{Z} with the unknowns r_1, \dots, r_t as follows: $\sum_{i=1}^t u_i \cdot r_i = \sum_{i=1}^t u_i \cdot x_i$. By collecting sufficiently many such vectors, we obtain r_1, \dots, r_t through solving those integer equations with the help of the bound of r_i and the LLL algorithm, then eventually can recover p via Euclidean algorithm. We actually find such a vector \mathbf{u} via the LLL reduction algorithm on the lattice \mathcal{L} which is generated with the row vectors of the matrix:

$$\begin{pmatrix} 1 & & & x_1 \\ & 1 & & x_2 \\ & & \ddots & \vdots \\ & & & 1 & x_{t-1} \\ & & & & x_t \end{pmatrix}.$$

The **coordinate vector** of the shortest vector of the LLL reduction gives us a vector $\mathbf{u} = (u_1, \dots, u_t)$ satisfying $\sum_{i=1}^t u_i \cdot r_i = \sum_{i=1}^t u_i \cdot x_i$. By computer experiments, from the LLL-reduced basis of lattice \mathcal{L} , we can find $t - 1$ such vectors, which allows us to find r_i and recover p via Euclidean algorithm.

We show that our algorithm is more efficiency than the algorithm proposed in [14] and we show that our algorithm can be used to attack the fully homomorphic encryption schemes, which are based on the approximate common divisors problem.

REFERENCES

- [1] A. K. Lenstra, H. W. Lenstra, Jr., and L. Lovász. Factoring polynomials with rational coefficients. *Mathematische Annalen*, 261:515-534, 1982.
- [2] Brakerski Z, Vaikuntanathan V. Efficient fully homomorphic encryption from (standard) LWE. *Foundations of Computer Science (FOCS)*, 2011 IEEE 52nd Annual Symposium on. IEEE, 2011: 97-106.

- [3] Brakerski Z. Fully homomorphic encryption without modulus switching from classical GapSVP. *Advances in Cryptology-CRYPTO 2012*. Springer Berlin Heidelberg, 2012: 868-886.
- [4] Bosma W, Cannon J, Playoust C. The Magma algebra system I: The user language. *Journal of Symbolic Computation*, 1997, 24(3): 235-265.
- [5] C. P. Schnorr. A more efficient algorithm for lattice basis reduction. *Journal of Algorithms*, 9(1):47-62, 1988.
- [6] Chen Y, Nguyen P Q. Faster algorithms for approximate common divisors: Breaking fully homomorphic encryption challenges over the integers. *Advances in Cryptology-EUROCRYPT 2012*. Springer Berlin Heidelberg, 2012: 502-519.
- [7] Cohn H, Heninger N. Approximate common divisors via lattices. *arXiv preprint arXiv:1108.2714*, 2011.
- [8] Coron J S, Mandal A, Naccache D, et al. Fully homomorphic encryption over the integers with shorter public keys. *Advances in Cryptology-CRYPTO 2011*. Springer Berlin Heidelberg, 2011: 487-504.
- [9] Coron J S, Naccache D, Tibouchi M. Public key compression and modulus switching for fully homomorphic encryption over the integers. *Advances in Cryptology-EUROCRYPT 2012*. Springer Berlin Heidelberg, 2012: 446-464.
- [10] Cheon J H, Coron J S, Kim J, et al. Batch fully homomorphic encryption over the integers. *Advances in Cryptology-EUROCRYPT 2013*. Springer Berlin Heidelberg, 2013: 315-335.
- [11] Coppersmith D. Finding a small root of a univariate modular equation. *Advances in Cryptology-EUROCRYPT96*. Springer Berlin Heidelberg, 1996: 155-165.
- [12] Gentry C. A fully homomorphic encryption scheme. Stanford University, 2009.
- [13] Gentry C, Halevi S. Implementing gentry's fully-homomorphic encryption scheme. *Advances in Cryptology-EUROCRYPT 2011*. Springer Berlin Heidelberg, 2011: 129-148.
- [14] Howgrave-Graham N. Approximate integer common divisors. *Cryptography and Lattices*. Springer Berlin Heidelberg, 2001: 51-66.
- [15] Stehlé D, Steinfeld R. Faster fully homomorphic encryption. *Advances in Cryptology-ASIACRYPT 2010*. Springer Berlin Heidelberg, 2010: 377-394.
- [16] Micciancio D, Goldwasser S. *Complexity of lattice problems: a cryptographic perspective*. Springer, 2002.
- [17] Novocin A, Stehlé D, Villard G. An LLL-reduction algorithm with quasi-linear time complexity. *Proceedings of the 43rd annual ACM symposium on Theory of computing*. ACM, 2011: 403-412.
- [18] Stehlé D, Zimmermann P. A binary recursive gcd algorithm. *Algorithmic number theory*. Springer Berlin Heidelberg, 2004: 411-425.
- [19] Van Dijk M, Gentry C, Halevi S, et al. Fully homomorphic encryption over the integers. *Advances in Cryptology-EUROCRYPT 2010*. Springer Berlin Heidelberg, 2010: 24-43.

Compact and tightly secure cryptography in the standard model

DENNIS HOFHEINZ

Tight security reductions. To argue for the security of a given cryptographic scheme S , we usually employ a security reduction. That is, we try to argue that every hypothetical adversary \mathcal{A}_S on S can be converted into an adversary \mathcal{A}_P on an allegedly hard computational problem P . In that sense, the only way to break S is to solve P . Of course, we are most interested in reductions to well-investigated problems P . Furthermore, there are reasons to consider the *tightness*

of the reduction: a tight reduction guarantees that \mathcal{A}_P 's success ε_P in solving P (in a reasonable metric) is about the same as \mathcal{A}_S 's success ε_S in attacking S .

To explain the impact of a (non-)tight reduction in more detail, consider a public-key encryption (PKE) scheme S that is deployed in a many-user environment. In this setting, an adversary \mathcal{A}_S on S may observe, say, n_C ciphertexts generated for each of the, say, n_U users. Most known security reductions in this setting are non-tight, in the sense that $\varepsilon_P \leq \frac{\varepsilon_S}{n_U \cdot n_C}$. As a consequence, keylength recommendations should also take n_U and n_C into account; no “universal” keylength recommendations can be given for such a scheme. This is particularly problematic in settings that grow significantly beyond initial expectations.

Tightly secure encryption and signature schemes. The construction of tightly secure cryptographic schemes appears to be a nontrivial task. For instance, although already explicitly considered in 2000, tightly secure PKE schemes have only been constructed very recently.¹ Moreover, the existing schemes have rather large ciphertexts, or require large parameters.

The situation for tightly secure signature schemes is somewhat brighter, but results are still limited. There are efficient signature schemes that are tightly secure under “ q -type” or interactive assumptions, or in the random oracle model. There are also more recent and somewhat less efficient schemes tightly secure under standard assumptions. Some of these latter schemes can even be converted into tightly secure PKE schemes; however, all of the resulting schemes suffer from asymptotically large parameters, keys, or signatures (resp. ciphertexts).

One difficulty in achieving tight security is that a tight reduction cannot afford to modify many challenges one by one. For instance, a common strategy to prove ciphertext indistinguishability (IND-CCA security) of a PKE scheme is to fully randomize all challenge ciphertexts given to an adversary. If this randomization is done via a hybrid argument over all challenge ciphertexts, randomizing one ciphertext at a time, then the reduction becomes non-tight. On the other hand, most existing strategies to construct efficient IND-CCA secure PKE schemes (such as hash proof systems, lossy trapdoor functions, or more specific strategies) are tailored towards randomizing only a single challenge ciphertext at a time.

Our contribution. We describe the first (almost) tightly secure signature and PKE schemes that are compact, in the sense that parameters, keys, and signatures (resp. ciphertexts) only contain a constant number of group elements. Our security reduction loses only a factor of $\mathbf{O}(k)$, where k is the security parameter. In particular, our security reduction does not degrade in the number of users or signatures, resp. ciphertexts. The security of our schemes is based upon the Decisional Diffie-Hellman (DDH) assumption in pairing-friendly groups.

¹We note that certain earlier PKE schemes achieve at least a certain form of tight security under nonstandard, “ q -type” assumptions, or in the random oracle model.

Ring switching and Bootstrapping Fully Homomorphic Encryption

CHRIS PEIKERT

We describe a technique for homomorphically evaluating any desired R -linear function $L : R' \rightarrow R$ on a ciphertext over R' (yielding a ciphertext over R), where R, R' are arbitrary cyclotomic rings and R is a subring of R' . The security of the method relies on the hardness of the ring-LWE problem over the base ring R .

We then extend the above method to give a quasi-linear $\tilde{O}(\lambda)$ algorithm for ‘bootstrapping’ a somewhat (for ‘packed’ ciphertexts over a ring) into a fully homomorphic one. The main technique involves switching through a sequence of ‘hybrid’ rings, which has the effect of homomorphically evaluating a discrete Fourier transform on a plaintext.

Physical zero knowledge

MONI NAOR

(joint work with Ben Fisch, Daniel Freund)

Is it possible to prove that two DNA-fingerprints match, or that they do not match, without revealing *any* further information about the fingerprints? Is it possible to prove that two objects have the same design without revealing the design itself? In the digital domain, *zero-knowledge* is an established concept where a prover convinces a verifier of a statement without revealing any information beyond the statement’s validity. However, zero-knowledge is not as well-developed in the context of problems that are *inherently physical*. In this paper, we are interested in protocols that prove physical properties of physical objects without revealing further information. The literature lacks a unified formal framework for designing and analyzing such protocols. We suggest the first paradigm for formally defining, modeling, and analyzing *physical zero-knowledge* (PhysicalZK) protocols, using the Universal Composability framework. We also demonstrate applications of physical zero-knowledge to DNA profiling and neutron radiography. Finally, we explore *public observation proofs*, an analog of public-coin proofs in the context of PhysicalZK.

Zero-knowledge proofs are protocols that prove an assertion without revealing any information beyond that assertion’s validity. Zero-knowledge proofs were first introduced by Goldwasser, Micali, and Rackoff in 1985. The power of zero-knowledge proofs is quite remarkable: anything that can be proved efficiently can be proved with a zero-knowledge protocol, under the cryptographic assumption that one-way functions exist (see Goldreich).

Zero-knowledge proofs have also been considered in a physical setting. A number of works have explored constructions of zero-knowledge protocols that can be physically implemented. One goal of those works was to design protocols with simple procedures and security arguments that the participating parties could easily understand. An added advantage of simple physical protocols is that humans can implement them without the aid of computers. Moran and Naor give methods

for polling people on sensitive issues using physical envelopes as an alternative to electronic polling, where humans might not trust computers to behave honestly. Many works have also addressed the incorporation of physical hardware into broader cryptographic schemes. In some cases, these hybrid protocols achieve efficiency or security gains that are unachievable in a standard computation model. Examples of physically realizable functionalities that have been suggested for aiding general cryptographic protocols include tamper-evidence, tamper-proof tokens, one-time programs, and physically uncloneable functions.

Previous literature on zero-knowledge in a physical setting addressed physical protocols for tasks that could otherwise be solved digitally. There is comparatively little formal work on protocols for inherently physical tasks that cannot be solved digitally. One example that has been studied rigorously is *distance bounding protocols*, introduced by Brands and Chaum in 1993, in which a verifier party determines or verifies an upper bound on its physical distance to a prover party. In 2012, Glaser, Barak, and Goldston suggested applying zero-knowledge concepts to the task of proving that a nuclear weapon is authentic without revealing sensitive information about its actual design, a problem that arises in the context of nuclear disengagement treaties. They presented an ϵ -knowledge protocol for this task, but did not have a rigorous framework for formally defining and analyzing the protocol's ϵ -knowledge security.

Our contributions. We present the first formal treatment of *physical zero-knowledge* (PhysicalZK) proofs for inherently physical claims. In our setting, a prover convinces a verifier that an input object satisfies a given physical property. Our framework for designing and analyzing PhysicalZK protocols uses the *Universally Composable* (UC) security framework, popularly applied in analysis of hybrid protocols involving physical hardware.

Expanding on Glaser et al., we present the first PhysicalZK protocols for the warhead verification problem, or the general task of verifying object neutron radiograph equality. Another application of PhysicalZK proofs is for DNA profiling in which a prover (e.g. a suspect) convinces a verifier (e.g. the police) that its DNA profile does not match a target profile (e.g. obtained from a crime scene) without revealing to the verifier any further information about the profiles, and discuss a protocol for parental testing. In particular we adapt the Goldwasser-Sipser set lower bound in order to obtain publicly observable protocols.

A further goal of our work is to initiate a rigorous study into the foundations of physical zero-knowledge. We point out both differences and similarities between physical and standard ZK where they arise. In particular, Section 3 compares the UC properties of physical vs. digital ZK, and Section 6 explores a physical analog of public coin proofs.

Learning with Errors in the Exponent

ÖZGÜR DAGDELEN

(joint work with Sebastian Gajek, Florian Göpfert)

We initiate the study of a novel class of group-theoretic intractability problems. Inspired by the theory of learning in presence of errors [1] we ask if noise in the exponent amplifies intractability. We put forth the notion of *Learning with Errors in the Exponent (LWEE)* and rather surprisingly show that various attractive properties known to exclusively hold for lattices carry over. Most notably are worst-case hardness and post-quantum resistance. In fact, LWEE's duality is due to the reducibility to two seemingly unrelated assumptions: learning with errors and the representation problem [2] in finite groups. For suitable parameter choices LWEE superposes properties from each individual intractability problem. The argument holds in the classical and quantum model of computation. We give the very first construction of a semantically secure public-key encryption system in the standard model. The heart of our construction is an "error recovery" technique inspired by [3] to handle critical propagations of noise terms in the exponent.

BLENDING GROUP AND LATTICE THEORY. The LWEE assumption reconciles the group theoretic structure of discrete log related problems with the algebraic simplicity of lattice theory. The technical idea behind the LWEE assumption can be summarized as planting an LWE sample $(\vec{a}, b = \langle \vec{a}, \vec{s} \rangle + e) \in \mathbb{Z}_q^n \times \mathbb{Z}_q$ in the exponent of a generator g of some group \mathcal{G} of order q . That is, the LWEE distribution consists of samples $(g^{\vec{a}}, g^{\langle \vec{a}, \vec{s} \rangle + e}) \in \mathcal{G}^n \times \mathcal{G}$ where \vec{a} is sampled uniformly from \mathbb{Z}_q^n , and $\vec{s} \leftarrow_R \chi^n$, $e \leftarrow_R \chi$ from some distribution χ . Similar to LWE, learning with errors in the exponent comes in two versions: The search version of LWEE asks to compute the secret vector \vec{s} while in the decisional variant one is supposed to distinguish $g^{\langle \vec{a}, \vec{s} \rangle + e}$ from a randomly sampled group element in \mathcal{G} .

EXISTENTIAL RELATIONS. In an attempt to confine LWEE, we prove that learning with errors in the exponent may take over the hardness from both theories. While striving for the existential relation to the family of group-theoretic assumptions, we infer a rather surprising connection to the (search) representation problem (ℓ -SRP) introduced by Brands [2]. We give a tight reduction from ℓ -SRP to the search version of the LWEE problem.

Looking at the decisional learning with errors in the exponent problem, we first introduce the decisional pendant of the representation problem (ℓ -DRP): Given a tuple $g, g_1, \dots, g_\ell, g^{x_1}, \dots, g^{x_\ell}, h$ from \mathcal{G} , where $x_1, \dots, x_\ell \leftarrow \chi$ are sampled from some distribution χ , one cannot distinguish between $\prod_{i=1}^{\ell} g_i^{x_i} = h$ and a randomly sampled value h in \mathcal{G} . In the same vain as done for the k-linear assumption, we show that ℓ -DRP becomes progressively harder to solve in generic group model. We then give a reduction from DRP to LWEE. We note that both of our reductions from the RP problem are tight, and they hold for (necessarily non-uniform) distributions χ , if the underlying RP holds for the representation sampled from the same distribution.

Investigating the relation to lattices, we show that an algorithm solving either the search or decisional LWEE problem efficiently can be turned into a successful attacker against the search or decisional LWE problem. Our reductions are tight and hold for (necessarily non-uniform) distributions χ as well.

A CONCRETE CRYPTOSYSTEM. In the light of LWEE we give a construction of a public-key encryption scheme. One may scale the magnitude to which the RP and LWE intractability contributes to the overall security of the system. The selection of parameters (e.g., modulus, dimension) offers a flexibility to fine-tune the cryptosystem's resilience against progress in attacking the underlying RP or LWE problem or the evolution of quantum computers. Concretely, one may choose to make the scheme short, post-quantum secure, or double-hard.

REFERENCES

- [1] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In *STOC*, pages 84–93, 2005.
- [2] Stefan A. Brands. An efficient off-line electronic cash system based on the representation problem. Technical report, Amsterdam, The Netherlands, The Netherlands, 1993.
- [3] Marc Joye and Benot Libert. Efficient cryptosystems from 2^k -th power residue symbols. *EUROCRYPT 2013*, volume 7881 of *LNCS*, pages 76–92. Springer Berlin Heidelberg, 2013.

An Almost-Optimally Fair Three-Party Coin-Flipping Protocol

IFTACH HAITNER

(joint work with Eliad Tzfatia)

In a multiparty *fair* coin-flipping (-tossing) protocol, the parties output a common (close to) unbiased bit, even though some corrupted parties try to bias the output. More formally, such protocols should satisfy the following two properties: first, when all parties are honest (i.e., follow the prescribed protocol), they all output the *same* bit, and this bit is unbiased (i.e., uniform over $\{0, 1\}$). Second, even when some parties are corrupted (i.e., collude and arbitrarily deviate from the protocol), the remaining parties should still output the *same* bit, and this bit should not be too biased (i.e., its distribution should be close to uniform over $\{0, 1\}$). We emphasize that, unlike weaker variants of coin-flipping protocol known in the literature, the honest parties should output a common bit, regardless of what the corrupted parties do. In particular, they are not allowed to abort if a cheat was noticed.

When a majority of the parties are honest, efficient and *completely* fair coin-flipping protocols are known as a special case of secure multiparty computation with an honest majority [BGW88].¹ When an honest majority is not guaranteed, however, the situation is more complex.

¹Throughout, we assume a broadcast channel is available to the parties.

Negative results. [Cle86] showed that for *any* efficient two-party m -round coin-flipping protocol, there exists an efficient adversary to bias the output of the honest party by $\Theta(1/m)$. This lower bound extends to the multiparty case via a simple reduction.

Positive results. Assuming one-way functions exist, [Cle86] showed that a simple m -round majority protocol can be used to derive a t -party coin-flipping protocol with bias $\Theta(\frac{\ell}{\sqrt{m}})$ (against dishonest majority), where ℓ is the number of corrupted parties. For more than two decades, [Cle86]’s protocol was the best known fair coin-flipping protocol (without honest majority), under *any* hardness assumption, and for *any* number of parties. In a recent breakthrough result, [MNS09] constructed an m -round, *two*-party coin-flipping protocol with optimal bias of $\Theta(\frac{1}{m})$. The result holds for any efficiently computable m , and under the assumption that oblivious transfer protocols exist. In a subsequent work, [BOO10] extended the result of [MNS09] for the multiparty case in which *less than* $\frac{2}{3}$ of the parties can be corrupted. More specifically, for any $\ell < \frac{2}{3} \cdot t$, they presented an m -round, t -party protocol, with bias $\frac{2^{\ell-t}}{m}$ against (up to) ℓ corrupted parties.

Still for the case of $\frac{2}{3}$ (or more) corrupted parties, the best known protocol was the $\Theta(\frac{\ell}{\sqrt{m}})$ -bias majority protocol of [Cle86]. In particular, this was the state of affairs for the natural three-party case (where two parties are corrupt).

Our result. We present an almost-optimally fair, three-party coin-flipping protocol. Specifically, assuming the existence of oblivious transfer protocols, we show that for any $m \in \text{poly}$ there exists an m -round, three-party coin-flipping protocol, with bias $\frac{O(\log^2 m)}{m}$ (against one, or two, corrupted parties).

As a building block towards constructing our three-party protocol, we present an alternative construction for two-party, almost-optimally fair coin-flipping protocols. Our approach does not follow the “threshold round” paradigm used in [MNS09, BOO10], but rather is a variation of the aforementioned $\Theta(\frac{\ell}{\sqrt{m}})$ -bias, coin-flipping protocol of [Cle86].

Open Problems. The existence of an optimally fair three-party coin-flipping protocol (without the $O(\log^2 m)$ factor) is still an interesting open question. A more fundamental question is whether there exists a fair coin-flipping protocol for any number of parties (against any number of corrupted parties). While constructing (at least, almost) optimally fair, m -round coin-flipping protocols for a constant (or even $\log(m)$) number of parties seems within the reach of our current technique, handling a super-logarithmic number of parties, not to mention $\Omega(m)$, seems to require a completely new approach, and may not be possible at all.

REFERENCES

- [BGW88] Michael Ben-Or, Shafi Goldwasser, and Avi Wigderson. Completeness theorems for non-cryptographic fault-tolerant distributed computation. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC)*, 1988.
- [BLOO11] Amos Beimel, Yehuda Lindell, Eran Omri, and Ilan Orlov. $1/p$ -secure multiparty computation without honest majority and the best of both worlds. pages 277–296, 2011.

- [Blu83] Manuel Blum. How to exchange (secret) keys. *ACM Transactions on Computer Systems*, 1983.
- [BOO10] Amos Beimel, Eran Omri, and Ilan Orlov. Protocols for multiparty coin toss with dishonest majority. In *Advances in Cryptology – CRYPTO 2010*, pages 538–557, 2010.
- [CI93] Richard Cleve and Russell Impagliazzo. Martingales, collective coin flipping and discrete control processes. Manuscript, 1993.
- [Cle86] Richard Cleve. Limits on the security of coin flips when half the processors are faulty. In *Proceedings of the 20th Annual ACM Symposium on Theory of Computing (STOC)*, pages 364–369, 1986.
- [MN05] Tal Moran and Moni Naor. Basing cryptographic protocols on tamper-evident seals. In *ICALP: Annual International Colloquium on Automata, Languages and Programming*, 2005.
- [MNS09] Tal Moran, Moni Naor, and Gil Segev. An optimally fair coin toss. In *Theory of Cryptography, Fifth Theory of Cryptography Conference, TCC 2009*, pages 1–18, 2009.

The Computational Benefit of Correlated Instances

HUIJIA LIN

(joint work with Irit Dinur, Shafi Goldwasser)

The starting point of this research is that instances of computational problems often do not exist in isolation. Rather, multiple and correlated instances of the same problem arise naturally in the real world. The *challenge* is how to gain computationally from instance correlations when they exist. We will be interested in settings where significant computational gain can be made in solving a single primary instance by having access to additional auxiliary instances which are correlated to the primary instance via the solution space.

We focus on Constraint Satisfaction Problems (CSPs), a very expressive class of computational problems that is well-studied both in terms of approximation algorithms and NP-hardness and in terms of average case hardness and usage for cryptography, e.g. Feige’s random 3-SAT hypothesis, Goldreich’s one way function proposal, learning-parity-with-noise, and others.

To model correlations between instances, we consider *generating processes* over search problems, where a primary instance I is first selected according to some distribution D (e.g. worst case, uniform, etc); then auxiliary instances I_1, \dots, I_T are generated so that their underlying solutions S_1, \dots, S_T each are a “perturbation” of a primary solution S for I . For example, S_t may be obtained by the probabilistic process of flipping each bit of S with a small constant probability.

We consider a variety of naturally occurring worst case and average case CSPs, and show how availability of a small number of auxiliary instances generated through a natural generating process, radically changes the complexity of solving the primary instance, from intractable to expected polynomial time. Indeed, at a high-level, knowing a logarithmic number of auxiliary instances enables a close polynomial time approximation of the primary solution, and when in addition the “difference vector” between the primary and the auxiliary solution is known, the

primary solution can be exactly found. Furthermore, knowing even a single auxiliary instance already enables finding the exact primary solution for a large class of CSPs.

An Algebraic Approach to Non-Malleability

ALON ROSEN

(joint work with Vipul Goyal, Silas Richelson, Margarita Vald)

In their seminal work on non-malleable cryptography, Dolev, Dwork and Naor, showed how to construct a non-malleable commitment with logarithmically-many "rounds"/"slots", the idea being that any adversary may successfully maul in some slots but would fail in at least one. Since then new ideas have been introduced, ultimately resulting in constant-round protocols based on any one-way function. Yet, in spite of this remarkable progress, each of the known constructions of non-malleable commitments leaves something to be desired.

We propose a new technique that allows us to to construct a non-malleable protocol with only a single "slot", and to improve in at least one aspect over each of the previously proposed protocols. Two direct byproducts of our new ideas are a four round non-malleable commitment and a four round non-malleable zero-knowledge argument, the latter matching the round complexity of the best known zero-knowledge argument (without the non-malleability requirement). The protocols are based on the existence of one-way permutations (or alternatively one-way functions with an extra round) and admit very efficient instantiations via standard homomorphic commitments and sigma protocols.

Theorem. *Assume the existence of a 2-round statistically binding commitment scheme (which holds if and only if one-way functions exist) then there is a 4-round non-malleable commitment scheme.*

Theorem. *Assume the existence of one-way functions. Then there is a 4-round black-box non-malleable zero-knowledge argument for every language in NP.*

Our analysis relies on algebraic reasoning, and makes use of error correcting codes in order to ensure that committers' tags differ in many coordinates. One way of viewing our construction is as a method for combining many atomic sub-protocols in a way that simultaneously amplifies soundness and non-malleability, thus requiring much weaker guarantees to begin with, and resulting in a protocol which is much trimmer in complexity compared to the existing ones.

The New Protocol. Suppose that committer C wishes to commit to message m , and let $t_1, \dots, t_n \in \mathbb{Z}$ be a sequence of tags that uniquely correspond to C 's identity (more on the tags later). Let **Com** be a statistically binding commitment scheme, and suppose that $m \in \mathbb{F}_q$ where $q > \max_i 2^{t_i}$. The protocol proceeds as follows:

- (1) C picks random $\mathbf{r} = (r_1, \dots, r_n) \in \mathbb{F}_q^n$ and sends **Com**(m), $\{\mathbf{Com}(r_i)\}_{i=1}^n$ to R ;

- (2) R sends C a query vector $\mathbf{alpha} = (\alpha_1, \dots, \alpha_n)$ where each α_i is drawn randomly from $[2^{t_i}] \subset \mathbb{F}_q$;
- (3) C sends R the response $\mathbf{a} = (a_1, \dots, a_n)$ where $a_i = r_i \alpha_i + m$;
- (4) C proves in ZK that the values \mathbf{a} (from step 3) are consistent with m and \mathbf{r} (from step 1).

The statistical binding property of the protocol follows directly from the binding of **Com**. The hiding property follows from the hiding of **Com**, the zero-knowledge property of the protocol used in step 4, and from the fact that for every i the receiver R observes only a single pair of the form (α_i, a_i) , where $a_i = r_i \alpha_i + m$.

Note the role of C 's tags in the protocol: t_i determines the size of the i -th coordinate's challenge space. Historically, non-malleable commitment schemes have used the tags as a way for the committer to encode its identity into the protocol as a mechanism to prevent M (whose tag is different from C 's tag) from mauling. In our protocol the tags play the same role, albeit rather passively. For example, though the size of the i -th challenge space depends on t_i , the size of the total challenge space depends only on the sum $\sum_{i=1}^n t_i$ of the tags. In particular, our scheme leaves open the possibility that the left and right challenge spaces might have the same size (in fact this will be ensured by our choice of tags). This raises a red flag, as previous works go to great lengths to set up imbalances between the left and right challenge spaces in order to force M to "give more information than it gets". Nevertheless, we are able to prove that any mauling attack will fail.

REFERENCES

- [1] Vipul Goyal, Silas Richelson, Alon Rosen, Margarita Vald: An Algebraic Approach to Non-Malleability. IACR Cryptology ePrint Archive 2014: 586 (2014). To appear in *FOCS 2014*.

The Fiat-Shamir Transformation in the Quantum Random Oracle Model

MARC FISCHLIN

(joint work with Özgür Dagdelen, Tommaso Gagliardini)

In the random oracle model (ROM) all protocol participants, including the adversary, get oracle access to a random function [5, 2]. This random function represents an idealized version of a public cryptographic hash function which displays no weaknesses, and which the parties can only use via its input/output behavior. The ROM facilitates the design of very efficient and provably secure protocols, although proofs in the ROM only provide heuristic security arguments in reality when the idealized hash function is eventually implemented by some concrete function.

The ROM is more and more now also applied in settings where the adversary may have quantum power. As pointed out in [3] these quantum capabilities of the adversary open up another attack strategy which is still compliant with the idea that the adversary does not exploit the inner structure of the hash function. Namely, the quantum adversary may now evaluate the concrete hash function in

superposition. To capture such attacks in the model, Boneh et al. [3] introduced the quantum random oracle model (QROM) in which the adversary can now ask the random function oracle about quantum states and receives, appropriately encoded, all answers in superpositions. This corresponds to the adversary's ability to evaluate the actual hash function on a quantum machine, without exploiting any structural properties beyond this purely technological advantage. Boneh et al. [3] show that switching from the ROM to the QROM can indeed make cryptographic protocols insecure.

One of the classical applications of the ROM is to turn interactive identification protocols between a prover and a verifier into (non-interactive) signature schemes via the Fiat-Shamir transformation [5]. If applied correctly, this transformation yields a secure signature scheme in the ROM [7]. Here we investigate the question if this security is also preserved in the QROM. If so, then using quantum-resistant cryptographic primitives for the identification protocol would potentially also give quantum-resistant signature schemes assuming idealized hash functions. Our main result, however, is negative in this regard: Basically, if the identification protocol is secure against so-called active attacks and the prover's first message in the identification protocol is independent of the secret key, then giving a security proof for the signature scheme seems to be hard [4]. Another negative result in this vein appears in a recent work by Ambainis et al. [1].

Nonetheless, we can provide some positive results, saying that the Fiat-Shamir transformation also works in the QROM if the first message in the identification protocol could have been generated by the (honest) verifier obviously, in such a way that only the prover could compute a matching randomness via some trapdoor information. We use this result to conclude that a (modification of a) signature scheme by Lubashevsky based on lattices [6] is secure in the QROM.

REFERENCES

- [1] Andris Ambainis, Ansis Rosmanis, Dominique Unruh: *Quantum Attacks on Classical Proof Systems - The Hardness of Quantum Rewinding*. IACR Cryptology ePrint Archive 2014: 296 (2014)
- [2] Mihir Bellare and Phillip Rogaway. *Random oracles are practical: A paradigm for designing efficient protocols*. ACM Conference on Computer and Communications Security, pages 62–73, 1993. ACM Press.
- [3] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, Mark Zhandry: *Random Oracles in a Quantum World*. ASIACRYPT 2011, pages 41-69, 2011, Springer, Berlin, Germany.
- [4] Özgür Dagdelen, Marc Fischlin, Tommaso Gagliardoni: *The Fiat-Shamir Transformation in a Quantum World*. ASIACRYPT (2) 2013, pages 62-81, 2013, Springer, Berlin, Germany.
- [5] Amos Fiat and Adi Shamir. *How to prove yourself: Practical solutions to identification and signature problems*. Advances in Cryptology – CRYPTO volume 263 of Lecture Notes in Computer Science, pages 186–194, 1987. Springer, Berlin, Germany.
- [6] Vadim Lyubashevsky. *Lattice signatures without trapdoors*. Advances in Cryptology – EUROCRYPT 2012, volume 7237 of Lecture Notes in Computer Science, pages 738-755, 2012. Springer, Berlin, Germany.
- [7] David Pointcheval and Jacques Stern. *Security arguments for digital signatures and blind signatures*. Journal of Cryptology, 13(3):361–396, 2000.

Lattice Algorithms for Learning with Errors

STEVEN GALBRAITH

Lattice-based cryptography provides interesting new cryptographic functions that are good candidates for being secure even against quantum adversaries. One of the most versatile computational assumptions for lattices is the learning with errors (LWE) problem [8]. An important topic is to determine the best algorithms for this problem or special cases of it. We define a variant of it now that may be useful to obtain efficient cryptosystems for some applications.

Definition: Let q be an odd prime and $n, m \in \mathbb{N}$. Let $\mathbf{s} \in \{0, 1\}^n$ or $\{-1, 0, 1\}^n$ be secret (column vector) that is sampled uniformly at random. Let \mathbf{A} be an $m \times n$ matrix over \mathbb{Z}_q chosen uniformly at random. Let \mathbf{e} be a length m column vector with entries sampled from a fixed error distribution (i.e., a discrete Gaussian with small standard deviation compared with q , or perhaps $\{-1, 0, 1\}^m$). Let $\mathbf{b} \equiv \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q}$. The **binary-LWE distribution** is the distribution on pairs $(\mathbf{A}, \mathbf{b}) \in \mathbb{Z}_q^{m \times n} \times \mathbb{Z}_q^m$ produced by the above process.

The **Decisional-binary-LWE** problem is to distinguish the binary-LWE distribution from the uniform distribution. The **(Computational-)binary-LWE** problem on input (\mathbf{A}, \mathbf{b}) from the binary-LWE distribution is to compute the most likely solution (\mathbf{s}, \mathbf{e}) .

There are hardness results for binary-LWE due to Brakerski, Langlois, Peikert, Regev and Stehlé [4] and Micciancio and Peikert [7]. It is necessary to take larger values for n compared with the traditional LWE case. One can also consider ring variants. Binary-Ring-LWE with errors in $\{-1, 0, 1\}^m$ is essentially the same as the NTRU problem.

The asymptotically fastest attack on LWE is due to Blum-Kalai-Wasserman (BKW), and it runs in time $2^{O(n)}$. Albrecht, Faugère, Fitzpatrick and Perret [1] have customised the BKW algorithm for the case of binary-LWE. The BKW algorithm runs faster in this case, but the complexity is still $2^{O(n)}$. One drawback of BKW is that it needs many samples (officially it requires $2^{O(n)}$ samples, but it seems this can be relaxed in practice).

We consider lattice attacks for several reasons. One is the issue of the number of LWE samples above. Another is because lattice algorithms usually work much better in small dimensions than the analysis predicts. Certainly, BKW is not the best attack on LWE for the kind of problems suggested for practical cryptosystems, and it is not primarily how we decide parameters for given security levels.

Given an LWE instance (\mathbf{A}, \mathbf{b}) one considers the lattice $L = \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{v} \equiv \mathbf{A}\mathbf{s} \pmod{q} \text{ for } \mathbf{s} \in \mathbb{Z}^n\}$. To solve LWE we want to find a lattice point $\mathbf{y} \equiv \mathbf{A}\mathbf{s} \pmod{q}$ close to \mathbf{b} . The basic approach of lattice attacks is to first compute a basis matrix \mathbf{B} for the lattice L . Then run BKZ lattice basis reduction on \mathbf{B} . Finally, the close vector is found using the embedding technique or some kind of enumeration (see Liu and Nguyen [6]).

Attacks on binary LWE. We now restrict to the case $\mathbf{s} \in \{-1, 0, 1\}^n$ and \mathbf{e} sampled from a discrete Gaussian with standard deviation $\sigma > 1$ (our results are

most interesting when σ is moderately large). It is clear that the lattice attack does not exploit the fact that \mathbf{s} is binary (it only depends on the difference \mathbf{e} to the nearest lattice point $\mathbf{A}\mathbf{s} \pmod{q}$).

Hence we translate the LWE instance $\mathbf{b} = \mathbf{A}\mathbf{s} + \mathbf{e} \pmod{q}$, where \mathbf{s}, \mathbf{e} are short, to an $(n + m) \times m$ ISIS instance

$$(\mathbf{A} | \mathbf{I}_m) \begin{pmatrix} \mathbf{s} \\ \mathbf{e} \end{pmatrix} \equiv \mathbf{b} \pmod{q}.$$

However, note that $\|\mathbf{s}\|$ (which is roughly $\sqrt{n}/2$) is much smaller in general than $\|\mathbf{e}\|$ (which is roughly $\sqrt{m}\sigma$). Hence it is natural to try to balance the problem.

To solve this problem compute any vector $\mathbf{y} \in \mathbb{Z}^m$ (not necessarily small) such that $\mathbf{A}\mathbf{y} \equiv \mathbf{b} \pmod{q}$ and then find a lattice point in L close to \mathbf{y} where

$$L = \{\mathbf{v} \in \mathbb{Z}^m : \mathbf{A}\mathbf{v} \equiv 0 \pmod{q}\}.$$

If \mathbf{v} is close to \mathbf{y} then $\mathbf{s} = \mathbf{y} - \mathbf{v}$ is a short vector such that $\mathbf{A}\mathbf{s} \equiv \mathbf{b} \pmod{q}$.

Let \mathbf{B} be the basis matrix for L . We re-scale the problem by multiplying the first n rows of \mathbf{B}, \mathbf{v} and \mathbf{y} by σ . Then (renaming those quantities) we have

$$\mathbf{y} - \mathbf{v} = \begin{pmatrix} \sigma\mathbf{s} \\ \mathbf{e} \end{pmatrix}.$$

The vectors $\sigma\mathbf{s}$ and \mathbf{e} now have similar-sized entries. The effect on the lattice problem is this: The new error vector $\begin{pmatrix} \sigma\mathbf{s} \\ \mathbf{e} \end{pmatrix}$ has similar norm to the original vector, whereas the lattice volume is increased by σ^n . For details and security estimations see [3].

Conclusions. The theoretical hardness results for binary-LWE require increasing n to $n \log(q) \approx n \log(n)$ to achieve the same level of security for LWE and binary-LWE. Our experiments suggest that this is overkill and that $n \log(\log(n))$ is more than sufficient, however more is research needed to clarify this.

There are a number of open questions:

- Is it possible to prove binary LWE is hard for smaller values of n (in other words, give a tighter reduction of LWE to binary-LWE)?
- Find better ways to “amplify” LWE in the case where the number of samples is limited. Do BKW with fewer samples.
- Exploit the ring structure to get improved versions of these attacks.
- Consider cryptographic applications of binary-LWE in detail to determine if the benefits of using binary-LWE to implement schemes are stronger than the additional costs from using larger matrices.

REFERENCES

- [1] Martin R. Albrecht, Carlos Cid, Jean-Charles Faugère, Robert Fitzpatrick and Ludovic Perret, On the Complexity of the BKW Algorithm on LWE, to appear in *Designs, Codes and Cryptography*. Published online 19 July 2013.
- [2] Martin R. Albrecht, Robert Fitzpatrick, and Florian Göpfert, On the Efficacy of Solving LWE by Reduction to Unique-SVP, to appear in *proceedings of 2013 International Conference on Information Security and Cryptology*.
- [3] Shi Bai and Steven D. Galbraith, Lattice Decoding Attacks on Binary LWE, in W. Susilo and Y. Mu (eds.), *ACISP 2014*, Springer LNCS 8544 (2014) 322–337.

- [4] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev and Damien Stehlé, Classical hardness of learning with errors, in D. Boneh, T. Roughgarden and J. Feigenbaum (eds.), STOC 2013, ACM (2013) 575–584.
- [5] Richard Lindner and Chris Peikert, Better key sizes (and attacks) for LWE-based encryption, in A. Kiayias (ed.), CT-RSA 20 11, Springer LNCS 6558 (2011) 319–339.
- [6] Mingjie Liu and Phong Q. Nguyen, Solving BDD by Enumeration: An Update, in E. Dawson (ed.), CT-RSA 2013, Springer LNCS 7779 (2013) 293–309.
- [7] Daniele Micciancio and Chris Peikert, Hardness of SIS and LWE with Small Parameters, in R. Canetti and J. A. Garay (eds.), CRYPTO 2013, Springer LNCS 8042 (2013) 21–39.
- [8] Oded Regev, On lattices, learning with errors, random linear codes, and cryptography, in H. N. Gabow and R. Fagin (eds.), STOC 2005, ACM (2005) 84–93.

Interactive Proofs of Proximity: Delegating Computation in Sublinear Time

GUY ROTHBLUM

(joint work with Salil Vadhan, Avi Wigderson)

We initiate a study of interactive proofs with sublinear time verifiers. These can be used by a sublinear-time client to delegate computations to a powerful but untrusted prover/server.

As in the study of sublinear-time algorithms, randomness is essential. Following the literature on property testing, we seek proof systems where the verifier accepts inputs in a language, and rejects (with high probability) inputs that are *far* from the language. We call such a system an “interactive proof of proximity.” We explore the power of this model, and show general upper and lower bounds.

How to Delegate Computations: The Power of No-Signaling Proofs

Yael Tauman Kalai, Ron D. Rothblum

(joint work with Ran Raz)

1. DELEGATION FOR \mathcal{P}

We construct a 1-round delegation scheme (i.e., argument-system) for every language computable in time $t = t(n)$, where the running time of the prover is $\text{poly}(t)$ and the running time of the verifier is $n \cdot \text{polylog}(t)$. In particular, for every language in \mathcal{P} we obtain a delegation scheme with almost linear time verification. Our construction relies on the existence of a computational sub-exponentially secure private information retrieval (PIR) scheme.

The proof exploits a curious connection between the problem of *computation delegation* and the model of *multi-prover interactive proofs that are sound against no-signaling (cheating) strategies*, a model that was studied in the context of multi-prover interactive proofs with provers that share quantum entanglement, and is motivated by the physical principle that information cannot travel faster than light.

For any language computable in time $t = t(n)$, we construct a multi-prover interactive proof (MIP) that is sound against no-signaling strategies, where the running time of the provers is $\text{poly}(t)$, the number of provers is $\text{polylog}(t)$, and the running time of the verifier is $n \cdot \text{polylog}(t)$.

In particular, this shows that the class of languages that have polynomial-time MIPs that are sound against no-signaling strategies, is exactly EXP. Previously, this class was only known to contain PSPACE.

To convert our MIP into a 1-round delegation scheme, we use the method suggested by Aiello *et al.* [ABOR00], which makes use of a PIR scheme. This method lacked a proof of security. We prove that this method is secure assuming the underlying MIP is secure against no-signaling provers.

2. ARGUMENTS OF PROXIMITY

An interactive proof of proximity (IPP) is an interactive protocol in which a prover tries to convince a *sublinear-time* verifier that $x \in \mathcal{L}$. Since the verifier runs in sublinear-time, following the property testing literature, it is only required that the verifier rejects inputs that are far from \mathcal{L} . In a recent work, Rothblum *et al.* [RVW13] constructed an IPP for every language computable by a low depth circuit.

We study the computational analogue, where soundness is required to hold only against *computationally bounded* cheating provers, and refer to such protocols as *interactive arguments of proximity*.

We construct *one-round* arguments of proximity for *every language* computable in time t , where the running time of the verifier is $o(n) + \text{polylog}(t)$ and the running time of the prover is $\text{poly}(t)$. We obtain this result by following the paradigm of Kalai *et al.* [KRR13a]: First, we construct a multi-prover interactive proof of proximity (MIPP) that is sound against *no-signaling* strategies, which is a strong notion of soundness inspired by quantum physics and the principal that information cannot travel faster than light. Then we show how to convert any such MIPP into a one-round argument of proximity.

The parameters of our protocols are similar to those obtained by Rothblum *et al.* We also give a lower bound, showing that in both cases, these parameters are close to optimal.

Finally, we observe that any one-round argument of proximity immediately yields a one-round delegation scheme (without proximity) where the verifier runs in *linear* time.

REFERENCES

- [ABOR00] William Aiello, Sandeep Bhatt, Rafail Ostrovsky, and S. Raj. Rajagopalan. Fast verification of any remote procedure call: Short witness-indistinguishable one-round proofs for NP. In *ICALP: Annual International Colloquium on Automata, Languages and Programming*, 2000.
- [KRR13a] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. Delegation for bounded space. In *STOC*, pages 565–574, 2013.

- [KRR13b] Yael Tauman Kalai, Ran Raz, and Ron D. Rothblum. How to delegate computations: The power of no-signaling proofs. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:183, 2013.
- [RVW13] Guy N. Rothblum, Salil P. Vadhan, and Avi Wigderson. Interactive proofs of proximity: delegating computation in sublinear time. In *STOC*, pages 793–802, 2013.

Garbled RAM from One-way Functions

RAFAIL OSTROVSKY

(joint work with Sanjam Garg, Steve Lu, Alessandra Scafuro)

Consider the following concrete problem. A user U wants to store a large dataset D on an untrusted server S and to outsource to S the ability to execute programs P_1, P_2, \dots on inputs x_1, x_2, \dots , where the programs may modify D , such that S does not learn anything about the inputs or how the database is being accessed or modified, except the output of the programs. One of the simplest motivating examples to consider is that of binary search. In this example, the running time of insecure binary search is $T = \text{poly}(\log n)$ for $n = |D|$.

We can solve this problem, for example, using garbled circuits, a technique introduced by Yao [3]. Garbled circuits allow a user to convert a circuit C into a garbled version \tilde{C} and an input x into a garbled version \tilde{x} in such a way that \tilde{C} can be evaluated on \tilde{x} to reveal the output $C(x)$ but nothing else. Crucially, \tilde{C} can be used one time only. Continuing with our binary search example, we can transform the code of binary search into a circuit C that has the dataset D hardwired. Then we compute \tilde{C}, \tilde{x} using garbled circuits and the problem is solved. The running time of this solution is $\text{poly}(|D|, \kappa)$ where κ is the security parameter, for each query x .

The overhead of this solution is a severe drawback. Indeed, it is exponential in the running time T of the program. Moreover, for each new query x' , the user needs to send a new garbled circuit \tilde{C} . In general, garbled circuits are inadequate to solve any problem where the program that we want to garble runs in time polylogarithmic in the size of D . We need a scheme that instead of working with circuits (which needs the entire dataset hardwired), can interoperate with previously stored data directly and gives an overhead that is proportional only to the running time T of the program.

Prior works have proposed different models to address the issues of this potentially exponential gap. We work in the specific RAM model where there is a RAM (Random Access Memory) that stores the data and the program code, and a small, stateless CPU that reads and writes to RAM. Furthermore, we are interested in a non-interactive solution where the user U sends a single message for each program P that S wants to run on D .

In [2] Lu and Ostrovsky introduce *garbled RAM*. Their construction works as follows. The program P is decomposed into a sequence of T CPU steps and each CPU step is represented as a circuit. Each CPU step reads and (potentially) writes one bit of the RAM. The RAM stores the dataset D . Then the garbling works as

follow. First the RAM is garbled in some manner with a master key k and sent to the server S . Second, each CPU circuit is garbled, using a scheme for garbling circuit. To enable each CPU circuit to read and write from the garbled RAM and to communicate with the CPU circuit for the next step, all CPU circuits have hardwired within them the secret key k that was used to garble the RAM. Indeed, the garbled circuits communicate to each other by sending encryption of labels, that can be decrypted by the user using the garbled values stored in the RAM. During the computation, a location i of the RAM might be read and updated several times. Therefore, location i is garbled using the following information: the index i , the key k , and the time t in which the location was accessed last. Because timestamps are used in the garbling of the RAM, it also required a mechanism that allows the circuits to calculate the time when a location was last written.

An important property of the garbled RAM is that the main RAM, D , is garbled only once at the beginning, and for any new query or program, the user only needs to send one garbled circuit per CPU time step. This solution is based on the minimal assumption of the existence of one-way functions, and the running time and space complexity of the online garbled CPU steps is $T \cdot \text{poly}(\log(n), \kappa)$.

The above solution requires a circularity assumption (see [1] for a detailed discussion) due to a subtle technical issue that comes up in the security proof. On a very high level the issue is the following. At each step j of the computation a garbled circuit \tilde{C}_j encrypts both labels to evaluate the next garbled circuit \tilde{C}_{j+1} , using two PRF evaluations under key k . The evaluator will obtain only one PRF evaluation from the RAM, that allows to decrypt one label only. The security of \tilde{C}_{j+1} relies on the fact that only one label can be decrypted by the evaluator, therefore depends on the security of the encryption computed under a key derived from k . However, the security of the encryption scheme relies on the fact that the key k remains secret, and therefore depends on the security of the circuit \tilde{C}_{j+1} that has k hardwired. Technically, the problem is that, when arguing security of the encryption scheme, we need to show a reduction to an adversary that does not know k . However, to run this reduction, we need to compute the garbled circuits $\{\tilde{C}_l\}_{l \in T}$ and for that we need to know k .

To avoid this circularity assumption, [1] use the following property: the key used to encrypt the labels for circuit \tilde{C}_{j+1} is independent of the keys hardwired in any circuit \tilde{C}_l for $l > j+1$. More specifically, [1], propose two schemes that do not need the circularity assumption. The first solution used the following observation: the garbled circuits need only information to encrypt, while the garbled memory needs to store the information to decrypt. Therefore, instead of using a symmetric-key encryption scheme, let us use a public key encryption scheme, where the garbled circuits have only the public key hardwired (that will be known to the reduction), while each location of the garbled memory corresponds to a special secret key that allows to decrypt the ciphertexts computed by the garbled circuits. More precisely their idea can be built off of Identity Based Encryption (IBE), where the circuits have hardwired the master public key MPK, and each location i of the memory corresponds to a secret key for identity i, t (where t is the time i was accessed

last). This solution fixes the problem and keeps the same overhead of the original solution, *but* it requires IBE which is a stronger assumption than OWFs.

The other solution proposed to dynamically “revoke” keys as the circuits access memory locations from one time step to the next. More precisely, in the original construction the key k_i , used to encrypt the labels for a circuit that need to read location i , is computed as the PRF evaluation under secret key k of the location i and time t . Their idea is to revoke the point k_i in the PRF evaluation once this point has been evaluated. This construction solves the circularity problem, while maintaining the minimal assumption of OWF. The downside of this construction is that circuits have to *remember* the keys that have been revoked so far: at each step j of computation, \tilde{C}_j needs to obtain the list of $j - 1$ PRF points evaluated so far. If the CPU runs for T steps, the overall cost will be $T^2 \cdot \text{poly}(\log(n), \kappa)$, and with a recursive solution, can be reduced to $T \cdot \min(T, n^\epsilon) \cdot \text{poly}(\log(n), \kappa)$ for any $\epsilon > 0$.

It is left as an open problem in [1] to construct a garbled RAM based only on OWF (without the circularity assumption) with overhead $\text{poly}(\log(n), \kappa)$.

Our result. In this work we resolve this open problem. First, instead of garbling the RAM with a master key, we garble each RAM location i with an *freshly* sampled random key $k_i \in \{0, 1\}^\kappa$. Thus we have this huge key $K = k_1, \dots, k_n$ that is κn bits, that obviously cannot be hardwired in the garbled circuits.

Thus, as second step, we store the huge key in a binary tree data structure: we arrange the keys on the leaves and build up a tree where each internal node has a key that allows the decryption of its two children. Thus, starting from the key of the root, one can navigate the tree and reach any key k_i in $\log n$ steps only. Intuitively, the goal here is that each garbled circuit \tilde{C}_j needs to remember only the key of the root in order to reach any actual key k_i . As a circuit navigates the tree, it will update all the nodes visited using fresh keys. For this purpose, each garbled circuit is equipped with $\log n$ fresh keys used to update the nodes visited along the path towards the leaf in position i . Note that there is no need to remember at which time a key was updated or a node navigated, because the keys are read on the fly. The only information the circuits need to be synchronized about is the key of the root.

The circularity problem is solved for the following reason. First, each garbled circuit does not need to memorize any key except the root of the garbled tree. Second, at each step, the visited path is in some sense “revoked” as it is replaced with a fresh path. An interesting property of our solution is that it does not requires timestamps and does not require any mechanism to remember when a location was accessed last.

Our solution only requires $\text{poly}(\kappa, \log(n))$ overhead to store the garbed tree and $\text{poly}(\log(n))$ additional overhead in the running time of the CPU circuit, required to navigate the tree. Therefore it achieves the same overhead as the original solution of [2] but without the circularity assumption.

REFERENCES

- [1] Craig Gentry, Shai Halevi, Steve Lu, Rafail Ostrovsky, Mariana Raykova, and Daniel Wichs. Garbled RAM revisited. In *EUROCRYPT* 2014.
- [2] Steve Lu and Rafail Ostrovsky. How to garble RAM programs. In *EUROCRYPT* 2013.
- [3] Andrew Chi-Chih Yao. Protocols for secure computations (extended abstract). In *FOCS*, pages 160–164, 1982.

Lattices with symmetry, and the extended tensor algebra

HENDRIK W. LENSTRA, ALICE SILVERBERG

For large ranks, there is no good algorithm that decides whether a given lattice has an orthonormal basis. But when the lattice is given with enough symmetry, we can construct a provably deterministic polynomial time algorithm to accomplish this, based on an algorithm of Gentry and Szydło in §7 of [1]. In addition, we put the Gentry-Szydło algorithm into a mathematical framework, and show that it is part of a general theory of “lattices with symmetry”.

A *lattice* (or integral lattice) is a finitely generated abelian group L equipped with a positive definite symmetric bilinear map $\langle \cdot, \cdot \rangle : L \times L \rightarrow \mathbb{Z}$. An *isometry* of lattices $L \xrightarrow{\sim} L'$ is a group isomorphism respecting $\langle \cdot, \cdot \rangle$. Let G be a finite abelian group equipped with an element u of order 2. We write $\mathbb{Z}\langle G \rangle$ for the modified group ring $\mathbb{Z}[G]/(u + 1)$; if $G = S \sqcup uS$, then $\mathbb{Z}\langle G \rangle = \bigoplus_{\sigma \in S} \mathbb{Z} \cdot \sigma$. We define an involution $a = \sum_{\sigma \in G} a_{\sigma} \sigma \mapsto \bar{a} = \sum_{\sigma \in G} a_{\sigma} \sigma^{-1}$ of $\mathbb{Z}\langle G \rangle$, and a \mathbb{Z} -linear map $t : \mathbb{Z}\langle G \rangle \rightarrow \mathbb{Z}$ where $t(\sum_{\sigma \in G} a_{\sigma} \sigma) = a_1 - a_u$; here $a_{\sigma} \in \mathbb{Z}$.

By a G -lattice we mean a lattice L equipped with a group homomorphism from G to the group of self-isometries of L , written $\sigma \mapsto (x \mapsto \sigma x)$, such that $ux = -x$ (with $x \in L$). An example is $L = \mathbb{Z}\langle G \rangle$, with $\langle x, y \rangle = t(x\bar{y})$, which has S as an orthonormal basis.

The main result is a deterministic polynomial time algorithm that given G and a G -lattice L decides whether there is an isometry $\mathbb{Z}\langle G \rangle \xrightarrow{\sim} L$ that respects the G -action, and if so exhibits one. The techniques involve algorithmic algebraic number theory, analytic number theory, commutative algebra, and lattice basis reduction, along with the methods of Gentry and Szydło [1] that this work extends.

The algorithm starts off by testing whether L is an *invertible* G -lattice, i.e., whether the map $L \otimes_{\mathbb{Z}\langle G \rangle} \bar{L} \rightarrow \mathbb{Z}\langle G \rangle$, $x \otimes \bar{y} \mapsto \sum_{\sigma \in S} \langle x, \sigma y \rangle \sigma$, is an isomorphism; here \bar{L} is a G -lattice with an isometry $L \xrightarrow{\sim} \bar{L}$, $x \mapsto \bar{x}$ satisfying $\overline{\sigma x} = \sigma \bar{x}$. If L is invertible, one can define the ring $\Lambda = \bigoplus_{i \in \mathbb{Z}} L^{\otimes i}$ (the “extended tensor algebra”), where $L^{\otimes 0} = \mathbb{Z}\langle G \rangle$, $L^{\otimes i} = L \otimes L \otimes \cdots \otimes L$ (with i L ’s) for $i > 0$, and $L^{\otimes i} = \bar{L} \otimes \bar{L} \otimes \cdots \otimes \bar{L}$ (with $-i$ \bar{L} ’s) for $i < 0$, where \otimes always means $\otimes_{\mathbb{Z}\langle G \rangle}$. The ring Λ has several useful theoretical and algorithmic properties, and forms the natural habitat for the computational techniques proposed by Gentry and Szydło.

See [2] for an extended abstract and see [3] for details. In future work we will extend the theory to “CM-orders”.

REFERENCES

- [1] C. Gentry and M. Szydło, *Cryptanalysis of the revised NTRU signature scheme*, Advances in Cryptology—EUROCRYPT 2002, Lect. Notes in Comp. Sci. **2332**, Springer, Berlin, 2002, 299–320.
- [2] H. W. Lenstra and A. Silverberg, *Revisiting the Gentry-Szydło algorithm*, in Advances in Cryptology—CRYPTO 2014, Part I, Lect. Notes in Comp. Sci. **8616**, Springer, Berlin, 2014, 280–296.
- [3] H. W. Lenstra and A. Silverberg, *Lattices with symmetry*, in preparation.

The Locality of Searchable Symmetric Encryption

STEFANO TESSARO

(joint work with David Cash)

Searchable symmetric encryption (SSE) enables a client to encrypt an index of record/keyword pairs and later issue tokens allowing an untrusted server to retrieve the (identifiers of) all records matching a keyword. SSE aims to hide statistics about the index to the greatest extent possible while maintaining practical efficiency for large indexes like email repositories or personal document stores. These schemes employ only fast symmetric primitives and recent implementations have shown that, in contrast to most applications of advanced cryptography, cryptographic processing like encryption is not the bottleneck for scaling. Instead, lower-level issues dealing with memory layouts required by the schemes are the limiting factor for large indexes.

This work studies how the security definitions for SSE inherently hamper scaling for large indexes. It proves an unconditional lower bound on the trade-off between server storage space and the *spatial locality* of its accesses to the encrypted index during a search. At a high level, the bound says that, for an index with N pairs, any secure SSE must either pad the encrypted index to an impractical (super-linear, $\omega(N)$) size *or* perform searching in a very non-local way (with $\omega(1)$ contiguous accesses or by reading far more bits than is necessary). Either of these options is likely to incur a large slow-down over a properly designed plaintext searching system with an $O(N)$ -size index that can search with $O(1)$ contiguous accesses.

The issue of locality in SSE surfaced in recent works where implementations showed that the non-local use of external storage was a bottleneck preventing scaling to large indexes. The only works with a highly local access pattern generated very large (roughly $O(N^2)$) encrypted databases that also prevented scaling. This paper explains this dichotomy of padding versus spatial locality by proving it is an unavoidable consequence of the SSE security definition. As more cryptographic applications are developed for securely outsourcing large amounts of data (while maintaining either authenticity or secrecy), lower-level issues like locality may become more relevant. While in some contexts (like secure multiparty computation) it is clear that the entire input must be touched during computation, this work appears to be the first to study of the effect of security on locality in detail.

The lower bound suggests the question of a matching upper bound. We give a new scheme with an $O(N \log N)$ size encrypted index and $O(\log N)$ locality via a different padding strategy, which compares to a scheme with a $O(N^2)$ size encrypted index and $O(1)$ locality. This scheme may not be competitive with prior highly-optimized implementations, but it serves as intermediate point in the trade-off curve implied by the lower bound. The interesting question of closing the gap is left open.

A Surprising Application of Differential Privacy

CYNTHIA DWORK

(joint work with Vitaly Feldman, Moritz Hardt, Toniann Pitassi, Omer Reingold, Aaron Roth)

False discovery is widely felt to be a growing problem in scientific research, and the past few decades have seen a great deal of effort to understand and propose mitigations for this problem. These efforts range from the use of holdout sets for estimating generalization error in machine learning, through sophisticated cross-validation techniques, to the use of deep statistical methods for controlling the false discovery rate in multiple hypothesis testing. Nonetheless, the theory surrounding this body of work assumes a fixed collection of hypotheses to be tested, or learning task to be achieved, selected non-adaptively before the data are gathered, whereas science is by definition an adaptive process, in which data are shared and re-used, and hypotheses and new studies are generated on the basis of data exploration and previous outcomes.

Although not usually understood in these terms, Freedman's paradox exemplifies the dangers of adaptivity: an equation is fitted, variables with small t statistics are dropped, and the – adaptively generated – new equation is refitted, with famously misleading results. When the relationship between the dependent and explanatory variables is weak, or even non-existent, the procedure overfits, erroneously “finding” significant relationships. We may think of this as generating a hypothesis for which the actual database is not representative of the distribution from which the database was drawn.

Differential privacy is a definition of privacy tailored to privacy-preserving data analysis [DMNS06, Dwo06]. A rapidly growing literature contains highly accurate differentially private algorithms for a broad range of common computational tasks. Roughly speaking, differential privacy ensures that the probability of observing any outcome from an analysis is “essentially unchanged” by changing any single database element. Here the probability distribution is over randomness introduced by the algorithm. Differentially private algorithms avoid overfitting because they do not rely too heavily on individual database elements.

Using insights derived from differential privacy we address the issue of adaptivity. We show that differentially private data exploration makes it hard to *find* a computation on which the data set is not representative. More precisely, it ensures that interacting with the data set yields no appreciable advantage in finding

such a computation, beyond what can be done without interacting with the data set. Thus, differential privacy neutralizes the risks due to adaptivity. A deeper examination of this phenomenon shows that the same holds whenever the data analyst's *choice* of the next computation to be carried out, or hypothesis to be tested, does not reveal too much information about the database. We formalize this sufficient condition with *max information*, a cross between mutual information and min entropy. Differential privacy is known to control the max information.

A common practice in machine learning is to use a *holdout set* to estimate the quality of the result of an analysis. For example, if the analysis learns a classifier then the holdout may be used to estimate the generalization error. In abstract and very simple terms, we can think of the holdout set as providing an oracle that makes Yes/No pronouncements on the validity of the outcome. The theory of holdout sets under adaptive computations is not developed: when can a holdout set be reused? What is to be done if the holdout indicates the outcome is not valid? Similarly, what is to be done when cross-validation procedures indicate a problematic outcome? The traditional approach is to recruit new data points sampled from the same distribution, which can be prohibitive when gathering new data is costly. In particular, this approach precludes sharing data sets for follow-up studies.

Our work provides a method for safely reusing a holdout set a great many times without undermining its accuracy as a validity oracle, even when hypotheses and computations are chosen adaptively. Armed with this technique, the analyst is free to explore the data *ad libitum*, generating and evaluating hypotheses and choosing new algorithms, verifying results on the holdout, and backtracking as needed.

This is achieved by partitioning the data into a training set and a holdout. The analyst is given unfettered access to the training set. Once she has reached her (possibly tentative) conclusions, she may check them against the holdout set; however, interactions with the holdout set are carried out in a differentially private fashion. The analyst is free to repeat this procedure a great – but not arbitrarily – many times.

By viewing the training set as part of the analyst's program, this setup is equivalent to an "expanded" analyst who interacts in a differentially private fashion with a data set. To be specific, the expanded analyst, who now has the training set hard-coded into her program, interacts with a differentially private holdout set. Since differential privacy neutralizes the risks inherent in adaptivity, the holdout set remains representative, that is, it serves the role of a "fresh" holdout set every time it is used.

REFERENCES

- [DMNS06] Cynthia Dwork, Frank McSherry, Kobbi Nissim, and Adam Smith. Calibrating noise to sensitivity in private data analysis. In *TCC '06*, pages 265–284, 2006.
- [Dwo06] C. Dwork. Differential privacy. In *Proceedings of the 33rd International Colloquium on Automata, Languages and Programming (ICALP)(2)*, pages 1–12, 2006.

On Virtual Grey Box Obfuscation for General Circuits

OMER PANETH

(joint work with Nir Bitansky, Ran Canetti, Yael Tauman Kalai)

Program obfuscation, namely the ability to efficiently compile a given program into a functionally equivalent program that is “unintelligible”, is an intriguing concept. Starting with the work of Barak et al. [2], a number of measures of security for program obfuscation have been proposed. Let us briefly review three notions of interest.

The first, *virtual black box (VBB)* obfuscation [2], requires that having access to the obfuscated program is essentially the same as having access to the program only as black box. Concretely, representing programs as circuits, an obfuscator \mathcal{O} for a family of circuits is VBB if for any poly-time adversary \mathcal{A} , there exists a poly-time simulator S , such that for any circuit C from the family, and any predicate $\pi(\cdot)$, \mathcal{A} cannot learn $\pi(C)$ from $\mathcal{O}(C)$ with noticeably higher probability than S can, given only oracle access to C . While this VBB obfuscation is natural and expressive, Barak et al. [2] showed that VBB is, in general, unobtainable.

A weaker variant of VBB, called *virtual grey-box (VGB)* [1], allows the simulator to be *semi-bounded*; namely, it can be computationally unbounded, while still making only a polynomial number of oracle queries to the circuit C . While significantly weaker than VBB in general, VGB is still meaningful for circuits that are unlearnable even by semi-bounded learners. Furthermore, VGB obfuscators for circuits escape the general impossibility results that apply to VBB obfuscators.

A weaker notion yet, called *indistinguishability obfuscation (IO)* [2], allows the (now computationally unbounded) simulator to also make an unbounded number of queries to C . Equivalently, \mathcal{O} is an IO for a circuit collection if for any two circuits C_0 and C_1 in the collection, having the same size and functionality, $\mathcal{O}(C_0)$ and $\mathcal{O}(C_1)$ are indistinguishable.

While IO has some attractive properties, and important cryptographic applications, the security guarantees provided by IO seem significantly weaker than those provided by either VBB or VGB obfuscation.

On the algorithmic level, for many years we had candidate obfuscators only for very simple functions such as point functions and variants. The landscape has changed completely with the recent breakthrough work of [6], which proposed a candidate general-purpose obfuscation algorithm for all circuits. [6] show that their scheme resists some simple attacks; but beyond that, they do not provide any analytic evidence for security.

Subsequently, considerable efforts have been made to analyze the security of the obfuscator in [6], and variants. The difficulty appears to be in capturing the security properties required from the *graded encodings schemes* [5], which is a central component in the construction. As a first step towards understanding the security of the obfuscator in [6], [4, 3] consider an ideal algebraic model, where the adversary is given “generic graded encodings” that can only be manipulated via admissible algebraic operations. They show that, in this model, variants of

the scheme in [6] are VBB obfuscators for all poly-size circuits. Still, neither of these idealized constructions or their analyses have, in of themselves, any bearing on the security of obfuscation algorithms in the plain model.

Pass et al. [7] make the first step towards proving the security of a general obfuscation scheme based on some natural hardness assumption in the plain model. They define a *semantic security* property for graded encoding schemes, aimed at capturing what it means for a graded encoding scheme to “behave as an ideal multi-linear graded encoding oracle”. They then show that a specially-crafted variant of the obfuscator of [3], with the ideal graded encoding scheme replaced by a semantically-secure graded encoding scheme, is IO for all circuits. *But what about stronger security notions?*

In this work, we obtain worst-case VGB obfuscation for NC^1 , based on a slight strengthening of the assumptions used in [7] to show IO for NC^1 . As an intermediate step towards this goal, we put forth a somewhat stronger variant of indistinguishability obfuscation, called *strong IO* (SIO). Informally, an obfuscator \mathcal{O} is SIO for a class of circuits \mathcal{C} if $\mathcal{O}(C) \approx \mathcal{O}(C')$ not only when $C, C' \in \mathcal{C}$ have the same functionality, but also when C and C' come from distributions over circuits in \mathcal{C} that are “close together”, in the sense that for any given input x , the probability that $C(x) \neq C'(x)$ is negligible.

We show that SIO is in fact *equivalent* to worst-case VGB obfuscation. Furthermore, for certain classes of functions, such as point functions, hyperplanes, or fuzzy point functions, SIO is equivalent to full-fledged worst-case VBB obfuscation. These equivalences hold unconditionally. Then, assuming existence of graded encoding schemes that satisfy a somewhat stronger variant of the semantic security notion of Pass et al. [7], we show that known obfuscation schemes are SIO for all circuits in NC^1 .

We also give evidence for the *necessity* of semantically-secure graded encoding for obtaining VGB. Specifically we show that, assuming existence of VGB obfuscators for all circuits, there exist *multilinear jigsaw puzzles* satisfying a form of semantic security. Multilinear jigsaw puzzles, defined in [6], are a limited-functionality variant of multilinear maps. They suffice for obtaining the positive result described above.

REFERENCES

- [1] Nir Bitansky and Ran Canetti. On strong simulation and composable point obfuscation. In *CRYPTO*, pages 520–537, 2010.
- [2] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *CRYPTO*, pages 1–18, 2001.
- [3] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. Cryptology ePrint Archive, Report 2013/631, 2013.
- [4] Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. Cryptology ePrint Archive, Report 2013/563, 2013.
- [5] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *EUROCRYPT*, pages 1–17, 2013.

- [6] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *FOCS*, 2013.
- [7] Rafael Pass, Sidharth Telang, and Karn Seth. Obfuscation from semantically-secure multilinear encodings. Cryptology ePrint Archive, Report 2013/781, 2013.

Indistinguishability Obfuscation from Semantically-Secure Multilinear Encodings

RAFAEL PASS

(joint work with Karn Seth, Sidharth Telang)

The goal of *program obfuscation* is to “scramble” a computer program, hiding its implementation details (making it hard to “reverse-engineer”), while preserving the functionality (i.e., input/output behavior) of the program. Precisely defining what it means to “scramble” a program is non-trivial: on the one hand, we want a definition that can be plausibly satisfied, on the other hand, we want a definition that is useful for applications.

Hada [Had00] and Barak, Goldreich, Impagliazzo, Rudich, Sahai, Vadhan, and Yang [BGI+01] show that simulation-based notion such as *virtual black-box obfuscation (VBB)* [BGI+01]—which, roughly speaking, require that everything that can be learned from the code of the obfuscated program can be simulated using just black-box access to the functionality—run into strong impossibility results.

We here focus on the notion of *indistinguishability obfuscation*, first defined by Barak *et al.* [BGI+01] and explored by Garg, Gentry, Halevi, Raykova, Sahai, and Waters [GGH+13b]. Roughly speaking, this notion requires that obfuscations $\mathcal{O}(C_1)$ and $\mathcal{O}(C_2)$ of any two *equivalent* circuits C_1 and C_2 (i.e., whose outputs agree on all inputs) from some class \mathcal{C} are computationally indistinguishable. In a very recent breakthrough result, Garg, Gentry, Halevi, Raykova, Sahai, and Waters [GGH+13b] provided the first candidate constructions of indistinguishability obfuscators for all polynomial-size circuits, based on so-called *multilinear graded encodings* [GGH13a]—for which candidate constructions were recently discovered in the seminal work of Garg, Gentry and Halevi [GGH13a], and more recently, alternative constructions were provided by Coron, Lepoint and Tibouchi [CLT13].

But despite these amazing developments, the following question remains open:

Can the security of general-purpose indistinguishability obfuscators be reduced to some “natural” intractability assumption?

The principal goal of the current paper is to make progress toward addressing this question.

Note that while the construction of indistinguishability obfuscation of Garg *et al.* is based on *some* intractability assumption, the assumption is very tightly tied to their scheme—in essence, the assumption stipulates that their scheme is a secure indistinguishability obfuscator. The VBB constructions of Brakerski and Rothblum [BR14] and Barak *et al.* [BGK+13] give us more confidence in the plausible security of their obfuscators, in that they show that at least “generic” attacks

– that treat multilinear encoding as if they were “physical envelopes” on which multilinear operations can be performed – cannot be used to break security of the obfuscators. But at the same time, non-generic attacks against their scheme are known – since general-purpose VBB obfuscation is impossible. Thus, it is not clear to what extent security arguments in the generic multilinear encoding model should make us more confident that these constructions satisfy e.g., a notion of indistinguishability obfuscation. In this work, we initiate a study of this question.

We define a notion of semantic security of multilinear graded encoding schemes, which stipulates security of class of algebraic “decisional” assumptions: roughly speaking, we require that for every nuPPT distribution D over two *constant-length* sequences \vec{m}_0, \vec{m}_1 and auxiliary elements \vec{z} such that all arithmetic circuits (respecting the multilinear restrictions and ending with a zero-test) are *constant* with overwhelming probability over (\vec{m}_b, \vec{z}) , $b \in \{0, 1\}$, we have that encodings of \vec{m}_0, \vec{z} are computationally indistinguishable from encodings of \vec{m}_1, \vec{z} . Assuming the existence of semantically secure multilinear encodings and the LWE assumption, we demonstrate the existence of indistinguishability obfuscators for all polynomial-size circuits. We additionally show that if we assume subexponential hardness, then it suffices to consider a *single* (falsifiable) instance of semantical security (i.e., that semantical security holds w.r.t to a particular distribution D) to obtain the same result.

We rely on the beautiful candidate obfuscation constructions of Garg et al [GGH+13b], Brakerski and Rothblum [BR14] and Barak et al [BGK+13] that were proven secure only in idealized generic multilinear encoding models, and develop new techniques for demonstrating security in the standard model, based only on semantic security of multilinear encodings (which trivially holds in the generic multilinear encoding model).

We also investigate various ways of defining an “uber assumption” (i.e., a super-assumption) for multilinear encodings, and show that the perhaps most natural way of formalizing the assumption that “any algebraic decision assumption that holds in the generic model also holds against nuPPT attackers” is false.

REFERENCES

- [BGI+01] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil Vadhan, and Ke Yang. On the (im) possibility of obfuscating programs. In *Advances in Cryptology CRYPTO 2001*, pages 1–18. Springer, 2001.
- [BR14] Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. In *TCC*, pages 1–25, 2014.
- [BGK+13] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. Cryptology ePrint Archive, Report 2013/631, 2013.
- [CLT13] Jean-Sébastien Coron, Tancreède Lepoint, and Mehdi Tibouchi. Practical multilinear maps over the integers. In *Advances in Cryptology, CRYPTO 2013*, pages 476–493, 2013.
- [GGH13a] Sanjam Garg, Craig Gentry, and Shai Halevi. Candidate multilinear maps from ideal lattices. In *Advances in Cryptology–EUROCRYPT 2013*, pages 1–17. Springer, 2013.

- [GGH+13b] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. *Proc. of FOCS 2013*, 2013.
- [Had00] Satoshi Hada. Zero-knowledge and code obfuscation. In *Advances in Cryptology—ASIACRYPT 2000*, pages 443–457. Springer, 2000.

How to Avoid Obfuscation Using Witness PRFs

MARK ZHANDRY

The goal of program obfuscation in cryptography is to scramble a program with the intention of hiding embedded secrets. Recently, Garg et al. [1] gave the first candidate construction of a program obfuscator, which has sparked a flurry of research showing many exciting uses of obfuscation. Such uses include functional encryption [1], short signatures and deniable encryption [2], multiparty key exchange and traitor tracing [3], and much more [4, 5, 6, 7, 8, 9].

While these results are exciting, instantiating these schemes with current candidate obfuscators [1, 10, 11, 12, 13] has several drawbacks:

- First, these obfuscators only build obfuscation for *formulas*. Getting obfuscation for all circuits currently requires an expensive boosting step involving obfuscating the decryption algorithm for a fully homomorphic encryption scheme.
- Second, all of these constructions first convert the formula into a branching program that is either very long (in the case of [1, 10, 11, 12]) or very wide (in the case of [13]). Then, the branching program is encoded in a multilinear map. Long branching programs require a high level of multilinearity, and long or wide programs both require many group elements.

Our Results. In this work, we show that for several applications of obfuscation, a weaker primitive we call *witness pseudorandom functions* (witness PRFs) actually suffices. Informally, a witness PRF for an NP language L is a PRF F such that anyone with a valid witness that $x \in L$ can compute $F(x)$, but for all $x \notin L$, $F(x)$ is computationally hidden. More precisely, a witness PRF consists of the following three algorithms:

- $\text{Gen}(\lambda, L, n)$ takes as input (a description of) an NP language L and instance length n , and outputs a secret function key fk and public evaluation key ek .
- $F(\text{fk}, x)$ takes as input the function key fk , an instance $x \in \{0, 1\}^n$, and produces an output y
- $\text{Eval}(\text{ek}, x, w)$ takes the evaluation key ek , and instance x , and a witness w for x , and outputs $F(\text{fk}, x)$ if w is a valid witness, \perp otherwise.

For security, we require that for any $x \in \{0, 1\}^n \setminus L$, $F(\text{fk}, x)$ is pseudorandom, even given ek and polynomially many PRF queries to $F(\text{fk}, \cdot)$.

Witness PRFs are closely related to the concept of smooth projective hash functions, and can be seen as a generalization of constrained PRFs [14, 15, 16] to arbitrary NP languages. We first show how to replace obfuscation with witness

PRFs for certain applications. We then show how to build witness PRFs from multilinear maps. Our witness PRF is more efficient than current obfuscation candidates, with similar efficiency to existing witness encryption constructions. Moreover, our construction relies on very natural assumptions about the underlying maps. Below, we list our applications of witness PRFs:

- **Multiparty non-interactive key exchange without trusted setup.** The first such scheme is due to Boneh and Zhandry [3], which is built from indistinguishability obfuscation (iO) and pseudorandom generators (PRGs). We give a closely related construction, where the obfuscator is replaced with a witness PRF, and prove that security still holds.
- **Poly-many hardcore bits.** Bellare, Stepanovs, and Tessaro[17] construct a hardcore function of arbitrary output size for any one-way function. They require differing inputs obfuscation[18, 6, 7], which is a form of knowledge assumption for obfuscators. We show how to replace the obfuscator with a witness PRFs that satisfies an extractability notion of security.
- **Reusable Witness Encryption.** Garg, Gentry, Sahai, and Waters [19] define and build the first witness encryption scheme from multilinear maps. Later, Garg et al. [1] show that indistinguishability obfuscation implies witness encryption. We show that witness PRFs are actually sufficient. We also define a notion of reusability for witness encryption, and give the first construction satisfying this notion.
- **Rudich Secret Sharing for mNP.** Rudich secret sharing is a generalization of secret sharing to the case where the allowed sets are instances of a monotone NP (mNP) language, and an allowed set of shares plus the corresponding witness are sufficient for learning the secret. Komargodski, Naor, and Yogev [9] give the first construction for all of mNP using witness encryption. We give a related protocol using witness PRFs that is reusable.
- **Fully distributed broadcast encryption.** Boneh and Zhandry [3] observe that certain families of key exchange protocols give rise to distributed broadcast encryption, where users generate their own secret keys. However, the notion has some limitations, which we discuss. We put forward the notion of *fully distributed* broadcast encryption which sidesteps these limitations, and give a construction where secret keys, public keys, and ciphertexts are short.

Open Questions. In terms of functionality, witness PRFs can be seen as lying somewhere between witness encryption and obfuscation; understanding how witness encryption, witness PRFs, and obfuscation compare is an important goal. For example, for what other applications of obfuscation do witness PRFs actually suffice? Can witness PRFs be built generically from any witness encryption scheme? In terms of efficiency, witness PRFs are much closer to witness encryption — can witness PRFs be strengthened even further while maintaining the efficiency of witness encryption?

REFERENCES

- [1] Sanjam Garg, Craig Gentry, Shai Halevi, Mariana Raykova, Amit Sahai, and Brent Waters. Candidate indistinguishability obfuscation and functional encryption for all circuits. In *Proc. of FOCS*, 2013.
- [2] Amit Sahai and Brent Waters. How to Use Indistinguishability Obfuscation: Deniable Encryption, and More. In *Proc. of STOC*, 2014.
- [3] Dan Boneh and Mark Zhandry. Multiparty key exchange, efficient traitor tracing, and more from indistinguishability obfuscation. In *Proceedings of CRYPTO*, 2014. Full version available at the Cryptology ePrint Archive, <http://eprint.iacr.org/2013/642>.
- [4] Susan Hohenberger, Amit Sahai, and Brent Waters. Replacing a random oracle: Full domain hash from indistinguishability obfuscation. In *Proc. of EuroCrypt*, 2014.
- [5] Sanjam Garg, Craig Gentry, Shai Halevi, and Mariana Raykova. Two-round secure mpc from indistinguishability obfuscation. In *Theoretical Cryptography Conference (TCC)*, 2014.
- [6] Elette Boyle, Kai-Min Chung, and Rafael Pass. On extractability obfuscation. In *Theoretical Cryptography Conference (TCC)*, 2014.
- [7] Prabhanjan Ananth, Dan Boneh, Sanjam Garg, Amit Sahai, and Mark Zhandry. Differing-inputs obfuscation and applications. Cryptology ePrint Archive, Report 2013/689, 2013. <http://eprint.iacr.org/>.
- [8] Omkant Pandey, Manoj Prabhakaran, and Amit Sahai. Obfuscation-based non-black-box simulation and four message concurrent zero knowledge for np. Cryptology ePrint Archive, Report 2013/754, 2013. <http://eprint.iacr.org/>.
- [9] Ilan Komargodski, Moni Naor, and Eylon Yogev. Secret-sharing for np from indistinguishability obfuscation. Cryptology ePrint Archive, Report 2014/213, 2014. <http://eprint.iacr.org/>.
- [10] Zvika Brakerski and Guy N. Rothblum. Virtual black-box obfuscation for all circuits via generic graded encoding. Cryptology ePrint Archive, Report 2013/563, 2013. <http://eprint.iacr.org/>.
- [11] Boaz Barak, Sanjam Garg, Yael Tauman Kalai, Omer Paneth, and Amit Sahai. Protecting obfuscation against algebraic attacks. In *Proc. of EuroCrypt*, 2014.
- [12] Rafael Pass, Karn Seth, and Sidharth Telang. Indistinguishability obfuscation from semantically-secure multilinear encodings. Cryptology ePrint Archive, Report 2013/781, 2013. <http://eprint.iacr.org/>.
- [13] Prabhanjan Ananth, Divya Gupta, Yuval Ishai, and Amit Sahai. Optimizing obfuscation: Avoiding barrington’s theorem. Cryptology ePrint Archive, Report 2014/222, 2014. <http://eprint.iacr.org/>.
- [14] Dan Boneh and Brent Waters. Constrained Pseudorandom Functions and Their Applications. *Advances in Cryptology (AsiaCrypt)*, pages 1–23, 2013.
- [15] Aggelos Kiayias, Stavros Papadopoulos, Nikos Triandopoulos, and Thomas Zacharias. Delegatable pseudorandom functions and applications. In *Proceedings ACM CCS*, 2013.
- [16] Elette Boyle, Shafi Goldwasser, and Ioana Ivan. Functional signatures and pseudorandom functions. Cryptology ePrint Archive, Report 2013/401, 2013.
- [17] Mihir Bellare, Igors Stepanovs, and Stefano Tessaro. Poly-many hardcore bits for any one-way function. Cryptology ePrint Archive, Report 2013/873, 2013. <http://eprint.iacr.org/>.
- [18] Boaz Barak, Oded Goldreich, Russell Impagliazzo, Steven Rudich, Amit Sahai, Salil P. Vadhan, and Ke Yang. On the (im)possibility of obfuscating programs. In *CRYPTO*, pages 1–18, 2001.
- [19] Sanjam Garg, Craig Gentry, Amit Sahai, and Brent Waters. Witness encryption and its applications. In *Proc. of STOC*, 2013.

Outsourcing Private RAM Computation

DANIEL WICHS

(joint work with Craig Gentry, Shai Halevi, Mariana Raykova)

In this talk¹ we consider the challenge of *privately outsourcing* computation where a weak client wants to execute a program on a remote server while hiding from it the raw data to be used in the computation. Moreover, we want to ensure that: (1) The client should perform significantly less work than executing the program on his own, and (2) The server should not have to do much more work than executing the program.

One method of outsourcing computation relies on fully homomorphic encryption (FHE), where the client simply encrypts her input and decrypts the output, and the server computes the program on encrypted data. Unfortunately, this solution requires the server to translate the program into a circuit and therefore work as hard as the *circuit size* of the computation, which in general, can be much larger than the work needed to execute the program on a random-access machine (RAM).

In this talk we describe *reusable garbled RAM* schemes, which offer the first solution to private outsourcing of RAM computation, where the server's work is only proportional to the RAM run-time of the computation and the client's work is essentially independent of the complexity of the computation altogether. In addition, these protocols are *non-interactive* and have the structure of *reusable garbled RAM* schemes.

Garbled Computation. Garbled circuits allow a client to garble a circuit C and then an input x in such a way that a server can use these garbled values to compute $C(x)$ without learning anything more about x . Until recently, all such known schemes became insecure if the server ever got to see more than one garbled input per garbled circuit. In particular, such schemes are not very useful in the context of outsourcing computation, since the client would have to create a fresh garbled circuit for each computation and therefore perform work proportional to the circuit size. Last year Goldwasser et al. described the first *reusable* circuit-garbling scheme [GKP+13] where the client can garble a single circuit and then garble many inputs to that circuit without losing security. This allows private outsourcing of circuit computation where the client only needs to do a one-time pre-processing step to garble the circuit, at a cost proportional to the circuit size. Also recently, Lu and Ostrovsky introduced the notion of *garbled RAM* [LO13, GHL+14]. Similar to garbled circuits, the client can garble a RAM program P , and later garble an input x in such a way that a server can use these garbled values to compute $P(x)$ without learning anything more about x . The complexity of garbling a RAM program (client complexity), the size of the garbled RAM, and the complexity of evaluating a garbled RAM (server complexity) are all proportional to the RAM run-time of the program rather than its circuit size. Just like in Yao's circuits, the scheme is *not* reusable and becomes completely insecure if the server sees more than a single garbled input per garbled program. In other words,

¹This talk outlines the results published in [GHRW14].

the client has to garble a fresh program for every computation, which requires as much work as doing the computation and therefore does not offer any savings in the context of outsourcing. The above raises the natural question whether we can obtain a *reusable garbled RAM* achieving the best of both worlds. In a reusable garbled RAM scheme, the client can garble a program P once as a potentially expensive pre-processing step, and later outsource many arbitrary computations of this program to a server by efficiently garbling fresh inputs x_i . The server can evaluate the garbled program on each garbled input in time proportional to the RAM complexity of the program. Furthermore, we would also like to do this in a setting where the client initially garbles a large *persistent memory* (e.g., database) and the programs can read/write to this memory.

Our Solutions. We describe the first solutions to the above problem of *reusable garbled-RAM*. As our “*basic*” solution, we describe a protocol that works in the setting *without* persistent memory, and requires the client to perform an expensive one-time pre-processing step to garble the program. As our “*best-case*” solution, we describe a protocol that also works in the more complex setting involving persistent memory (e.g., database) and does not require any expensive pre-processing. Our solutions are built from *non-reusable garbled RAM* in conjunction with new types of *reusable garbled circuits* that are more efficient than prior solutions but only satisfy weaker security. For the basic setting without a persistent database, we can instantiate the required type of reusable garbled circuits from *indistinguishability obfuscation* or from *functional encryption for circuits* as a black-box. For the more complex setting with a persistent database, we can instantiate the required type of reusable garbled circuits using stronger notions of obfuscation. It remains an open problem to get solutions under weaker assumptions.

REFERENCES

- [GHRW14] Craig Gentry, Shai Halevi, Mariana Raykova, and Daniel Wichs. Outsourcing Private RAM Computation. FOCS, 2014.
- [GKP+13] Shafi Goldwasser, Yael Tauman Kalai, Raluca A. Popa, Vinod Vaikuntanathan, and Nikolai Zeldovich. Reusable garbled circuits and succinct functional encryption. STOC, 2013.
- [LO13] Steve Lu and Rafail Ostrovsky. How to garble RAM programs. EUROCRYPT, 2013.
- [GHL+14] Craig Gentry, Shai Halevi, Steve Lu, Rafail Ostrovsky, Mariana Raykova, and Daniel Wichs. Garbled RAM, revisited. EUROCRYPT, 2014.

Practical Multi-Party Computation

NIGEL P. SMART

In recent years actively secure MPC has moved from a theoretical subject into one which is becoming more practical. In the variants of multi-party computation which are based on secret sharing the major performance improvement has come from the technique of authenticating the shared data and/or the shares themselves using information theoretic message authentication codes (MACs). This idea has been used in a number of works: In the case of two-party MPC for binary circuits in

[10], for n -party dishonest majority MPC for arithmetic circuits over a “largish” finite field [3, 6], and for n -party dishonest majority MPC over binary circuits [7]. All of these protocols are in the pre-processing model, in which the parties first engage in a function and input independent offline phase. The offline phase produces various pieces of data, often Beaver style [2] “multiplication triples”, which are then consumed in the online phase when the function is determined and evaluated.

In the case of the protocol of [10], called Tiny-OT in what follows, the authors use the technique of applying information theoretic MACs to the oblivious transfer (OT) based GMW protocol [8] in the two party setting. In this protocol the offline phase consists of producing a set of pre-processed random OTs which have been authenticated. The offline phase is then executed efficiently using a variant of the OT extension protocol of [9]. For a detailed discussion on OT extension see [1, 9, 10]. In this talk we shall take OT extension as a given sub-procedure.

One can think of the Tiny-OT protocol as applying the authentication technique of [3] to the two party, binary circuit case, with a pre-processing which is based on OT as opposed to semi-homomorphic encryption. For two party protocols over binary circuits practical experiments show that Tiny-OT far out-performs other protocols, such as those based on Yao’s garbled circuit technique. This is because of the performance of the offline phase of the Tiny-OT protocol. Thus a natural question is to ask, whether one can extend the Tiny-OT protocol to the n -party setting for binary circuits.

In this talk we mainly address ourselves to the above question, i.e. how can we generalize the two-party protocol from [10] to the n -party setting?

We first describe what are the key technical difficulties we need to overcome. The Tiny-OT protocol at its heart has a method for authenticating random bits via pairwise MACs, which itself is based on an efficient protocol for OT-extension. In [10] this protocol is called aBit. Our aim is to use this efficient two-party process as a black-box. Unfortunately, if we extend this procedure naively to the three party case, we would obtain (for example) that parties P_1 and P_2 could execute the protocol so that P_1 obtains a random bit and a MAC, whilst P_2 obtains a key for the MAC used to authenticate the random bit. However, party P_3 obtains no authentication on the random bit obtained by P_1 , nor does it obtain any information as to the MAC or the key.

To overcome this difficulty, we present a protocol in which we fix an unknown global random key and where each party holds a share of this key. Then by executing the pairwise aBit protocol, we are able to obtain a secret shared value, as well as a shared MAC, by all n -parties. This resulting MAC is identical to the MAC used in the SPDZ protocol from [5]. This allows us to obtain authenticated random shares, and in addition to permit parties to enter their inputs into the MPC protocol.

The online phase will then follow similarly to [5], if we can realize a protocol to produce “multiplication triples”. In [10] one can obtain such triples by utilizing a complex method to produce authenticated random OTs and authenticated

random ANDs (called aOTs and aANDs)¹. We notice that our method for obtaining authenticated bits also enables us to obtain a form of authenticated OTs in a relatively trivial manner, and such authenticated OTs can be used directly to implement a multiplication gate in the online phase.

Our contribution is twofold. First, we generalize the two-party Tiny-OT protocol to the n -party setting, using a novel technique for authentication of secret shared bits, and completely new offline and online phases. Thus we are able to dispense with the protocols to generate aOTs and aANDs from [10], obtaining a simple and efficient online protocol. Second, and as a by product, we obtain a more efficient protocol than the original Tiny-OT protocol, in the two party setting when one measures efficiency in terms of the number of aBit's needed per multiplication gate. The security of our protocols are proven in the standard universal composability (UC) framework [4] against a malicious adversary and static corruption of parties.

REFERENCES

- [1] G. Asharov, Y. Lindell, T. Schneider, and M. Zohner. More efficient oblivious transfer and extensions for faster secure computation. In A.-R. Sadeghi, V. D. Gligor, and M. Yung, editors, *ACM Conference on Computer and Communications Security*, pages 535–548. ACM, 2013.
- [2] D. Beaver. Efficient multiparty protocols using circuit randomization. In J. Feigenbaum, editor, *CRYPTO*, volume 576 of *Lecture Notes in Computer Science*, pages 420–432. Springer, 1991.
- [3] R. Bendlin, I. Damgård, C. Orlandi, and S. Zakarias. Semi-homomorphic encryption and multiparty computation. In K. G. Paterson, editor, *EUROCRYPT*, volume 6632 of *Lecture Notes in Computer Science*, pages 169–188. Springer, 2011.
- [4] R. Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *FOCS*, pages 136–145. IEEE Computer Society, 2001.
- [5] I. Damgård, M. Keller, E. Larraia, V. Pastro, P. Scholl, and N. P. Smart. Practical covertly secure mpc for dishonest majority - or: Breaking the spdz limits. In J. Crampton, S. Jajodia, and K. Mayes, editors, *ESORICS*, volume 8134 of *Lecture Notes in Computer Science*, pages 1–18. Springer, 2013.
- [6] I. Damgård, V. Pastro, N. P. Smart, and S. Zakarias. Multiparty computation from somewhat homomorphic encryption. In Safavi-Naini and Canetti [11], pages 643–662.
- [7] I. Damgård and S. Zakarias. Constant-overhead secure computation of boolean circuits using preprocessing. In A. Sahai, editor, *TCC*, volume 7785 of *Lecture Notes in Computer Science*, pages 621–641. Springer, 2013.
- [8] O. Goldreich, S. Micali, and A. Wigderson. How to play any mental game or a completeness theorem for protocols with honest majority. In A. V. Aho, editor, *STOC*, pages 218–229. ACM, 1987.
- [9] Y. Ishai, J. Kilian, K. Nissim, and E. Petrank. Extending oblivious transfers efficiently. In D. Boneh, editor, *CRYPTO*, volume 2729 of *Lecture Notes in Computer Science*, pages 145–161. Springer, 2003.
- [10] J. B. Nielsen, P. S. Nordholt, C. Orlandi, and S. S. Burra. A new approach to practical active-secure two-party computation. In Safavi-Naini and Canetti [11], pages 681–700.

¹In fact the paper [10] does not produce such multiplication triples, but they follow immediately from the presentation in the paper and would result in a more efficient online phase than that described in [10]

- [11] R. Safavi-Naini and R. Canetti, editors. *Advances in Cryptology - CRYPTO 2012 - 32nd Annual Cryptology Conference, Santa Barbara, CA, USA, August 19-23, 2012. Proceedings*, volume 7417 of *Lecture Notes in Computer Science*. Springer, 2012.

Factoring Integers by CVP Algorithms

CLAUS P. SCHNORR

Let $N \in \mathbb{N}$ have distinct prime factors all greater than the first n primes p_1, \dots, p_n . We can factor N by solving **CVP**'s for the prime number lattice $\mathcal{L}(\mathbf{B}_{n,c})$ with basis matrix $\mathbf{B}_{n,c} = [\mathbf{b}_1, \dots, \mathbf{b}_n] \in \mathbb{R}^{(n+1) \times n}$ and the target vector $\mathbf{N} \in \mathbb{R}^{n+1}$:

$$\mathbf{B}_{n,c} = \begin{bmatrix} \sqrt{\ln p_1} & 0 & 0 \\ 0 & \ddots & 0 \\ 0 & 0 & \sqrt{\ln p_n} \\ N^c \ln p_1 & \cdots & N^c \ln p_n \end{bmatrix}, \quad \mathbf{N} = \begin{bmatrix} 0 \\ \vdots \\ 0 \\ N^c \ln N \end{bmatrix}.$$

We identify each $\sum_{i=1}^n u_i \mathbf{b}_i \in \mathcal{L}(\mathbf{B}_{n,c})$ with the pair of p_n -smooth $u = \prod_{u_i > 0} p_i^{u_i}$, $v = \prod_{u_i < 0} p_i^{-u_i} \in \mathbb{N}$. The goal of the **CVP** algorithm is to find u, v such that $|u - vN|$ is so small that it most likely is p_n -smooth and thus yields a non trivial relation $\prod_{i=1}^n p_i^{e_i} = \pm 1 \pmod N$ with $e_i \in \mathbb{Z}$. Given n such independent mod N -relations we can factor N . This **CVP** method generates mod N -relations given by p_n -smooth triples $u, v, |u - vN| \in \mathbb{N}$. Recent improvements:

- (1) We perform the stages in enumerating vectors $\mathbf{b} \in \mathcal{L}(\mathbf{B}_{n,c})$ close to \mathbf{N} in the order of their success rate to find some \mathbf{b} very close to \mathbf{N} . The success rate is based on the GAUSSIAN volume heuristics [4].
- (2) We extremely prune the enumeration of lattice vectors close to \mathbf{N} so that a very small, but sufficient fraction of these vectors gets efficiently generated.
- (3) We randomly multiply each of the first n rows of a BKZ-reduced basis of $\mathcal{L}(\mathbf{B}_{n,c})$ with probability $1/2$ by 2 before enumerating lattice vectors close to \mathbf{N} per round. This random scaling generates independent mod N -relations per round.
- (4) While this **CVP** method finds p_n -smooth triples $u, v, |u - vN|$ we must extend the method to generate relations with arbitrary $v \in \mathbb{N}$. This is because there do not exist enough mod N -relations with p_n -smooth v for very large N .

Under heuristic assumptions the **CVP** problem for the lattice $\mathcal{L}(\mathbf{B}_{n,c})$ and target vector \mathbf{N} can be solved in polynomial time [5, Cor.3]. This is because the

relative density $rd(\mathcal{L})$ of $\mathcal{L}(\mathbf{B}_{n,c})$ satisfies $rd(\mathcal{L}) = o(n^{-1/4})$ for large n . By definition $rd(\mathcal{L})$ transforms the Hermite inequality $\lambda_1 \leq \gamma_n^{1/2} \det(\mathcal{L})^{1/n}$ with Hermite constant γ_n into the equation $\lambda_1^2 = rd(\mathcal{L})^2 \gamma_n \det(\mathcal{L})^{2/n}$ and thus $0 < rd(\mathcal{L}) \leq 1$.

In order to solve our **CVP**'s in pol. time we need some nearly shortest lattice vector. The **SVP** to find $\mathbf{b} \in \mathcal{L}(\mathbf{B}_{n,c})$ of length λ_1 can be solved by extreme pruning under heuristics in pol. time [5, Prop.1] because $rd(\mathcal{L}) = o(n^{-1/4})$. Moreover, the analysis of extreme pruning of [1] with an heuristic success rate $1/n$ has been improved to a success rate $1 - o(1)$ in [2, chapter 4].

Right now we create one mod N -relations for $N \approx 10^{14}$ using $n = 90$ primes in 6 seconds per relation. There are $6.4 \cdot 10^5$ p_n -smooth triples $u, v, |u - vN| \leq p_n^3$.

For $N \gg 2^{14}$ there exist enough mod N -relations only for possibly non smooth $v \in \mathbb{N}$. Our method for directing and pruning the search towards successful v can be extended from p_n -smooth to arbitrary v : during the enumeration of the $\mathbf{b} \in \mathcal{L}(\mathbf{B}_{n,c})$ close to \mathbf{N} we iteratively replace N by vN and adjust the target vector \mathbf{N} to $\mathbf{N}_v = \mathbf{N} \ln vN / \ln N$ so that finally $\|\mathcal{L}(\mathbf{B}_{n,c}) - \mathbf{N}_v\|$ becomes so small that there exists a p_n -smooth u with $|u - vN| \leq p_n^3$.

Towards new record factorisations. For arbitrary $N \approx 2^{800}$ and $n = 900$ primes there exist 10^{11} p_n -smooth $u \leq N^3$ such that $|u - vN| \leq p_n^3$ holds for some $v \in \mathbb{N}$. This should enable the efficient generation of 900 mod N -relations.

We extend the proof of Lemma 5.3 of [3], we prove $\lambda_1^2 \geq 2c \ln N + 1 - N^{-c+c/n} \sqrt{n}$ for the lattice $\mathcal{L}(\mathbf{B}_{n,c})$ even when the prime 2 is in the prime basis [5, Le. 2].

A particular advantage of our factoring method is that the prime basis is much smaller than for all other known factoring methods.

REFERENCES

- [1] N. Gama, P.Q. Nguyen and O. Regev, Lattice enumeration using extreme pruning, Proc. EUROCRYPT 2010, LNCS 6110, Springer-Verlag, (2010). 257–278, final version to appear.
- [2] B. Lange, *Neue Schranken für SVP-Approximation und SVP-Aufzählungsalgorithmen*. Dissertation Frankfurt 2013 (D 30) //www.math.uni-frankfurt.de/~dmst/ Ph.D. Theses.
- [3] D. Micciancio and S. Goldwasser, *Complexity of Lattice Problems: A Cryptographic Perspective*. Kluwer Academic Publishers, Boston, London, (2002).
- [4] C.P. Schnorr and H.H. Hörner, *Attacking the Chor–Rivest cryptosystem by improved lattice reduction* In Proc. EUROCRYPT'95, LNCS 921, Springer-Verlag, (1995), 1–12. //www.mi.informatik.uni-frankfurt.de/
- [5] C.P. Schnorr, *Factoring integers by CVP Algorithms*, Proceedings Number Theory and Cryptography, LNCS 8260, Springer-Verlag, (2013), 73–93, for a more complete version see //www.mi.informatik.uni-frankfurt.de/

Non-Black-Box Simulation in the Fully Concurrent Setting

VIPUL GOYAL

Zero-knowledge proofs have played a central role in the design of secure cryptographic schemes. Introduced in [GMR89], all initial zero-knowledge (ZK) protocols shared a simple structure: the messages of the verifier in the protocol were simply random coin tosses. This simple structure is quite appealing in and by itself beyond any applications. However over a period of time, this public coin property found

applications in several (even seemingly unrelated) contexts. An (incomplete) list of such examples include: the Fiat-Shamir paradigm [FS86], zero-knowledge protocols for **IP** [BOGG88], efficient parallel repetition theorems [PV07], etc. Much of the early work on zero-knowledge was for the “stand-alone” setting where there is a single protocol execution running in isolation.

In a breakthrough work in 2001, Barak [Bar01] introduced non-black-box simulation techniques in cryptography. This was done by giving a protocol which was public-coin, constant rounds and secure with a bounded number of concurrent protocol executions. A key feature of the construction was that it did not rely on the traditional paradigm of rewinding the adversary (and instead the simulator was “straightline”). Barak’s technique has since then been utilized in a variety of different contexts and has been used to get results provable impossible using the traditional black-box simulation based on rewinding techniques.

The fact the protocol is secure only in the bounded concurrent setting has been an important limitation of Barak’s construction [Bar01]. There has been a long line of works on developing simulation strategies based on rewinding that would work in the fully concurrent setting (see [PRS02] and the references there in). However all these works rely on the paradigm of “extracting some trapdoor” from the adversarial verifier. This necessarily means that the protocol must not be public-coin. The existence of a public-coin concurrent zero-knowledge protocol has remained an intriguing question till now.

We present a new zero-knowledge argument protocol by relying on the non-black-box simulation technique of Barak [Bar01]. Similar to the protocol of Barak, ours is public-coin, is based on the existence of collision-resistant hash functions, and, is not based on “rewinding techniques” but rather uses non-black-box simulation. However in contrast to the protocol of Barak, our protocol is secure even if there are any unbounded (polynomial) number of concurrent sessions.

This gives us the first construction of public-coin concurrent zero-knowledge. Prior to our work, Pass, Tseng and Wikström [PTW11] had even shown that using black-box simulation, getting a construction for even public-coin parallel zero-knowledge is impossible.

A public-coin concurrent zero-knowledge protocol directly implies the existence of a concurrent resettably-sound zero-knowledge protocol. This is an improvement over the corresponding construction of Deng, Goyal and Sahai [DGS09] which was based on stronger assumptions. Furthermore, this also directly leads to an alternative (and arguable cleaner) construction of a simultaneous resettable zero-knowledge argument system.

An important feature of our protocol is the existence of a “straight-line” simulator. This gives a fundamentally different tool for constructing concurrently secure computation protocols (for functionalities even beyond zero-knowledge).

The round complexity of our protocol is n^ϵ (for any constant $\epsilon > 0$), and, the simulator runs in strict polynomial time. The main technique behind our construction is *purely combinatorial* in nature.

REFERENCES

- [GMR89] Shafi Goldwasser, Silvio Micali, and Charles Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.
- [FS86] Amos Fiat and Adi Shamir. How to prove yourself: Practical solutions to identification and signature problems. In *CRYPTO*, pages 186–194, 1986.
- [BOGG88] Michael Ben-Or, Oded Goldreich, Shafi Goldwasser, Johan Håstad, Joe Kilian, Silvio Micali, and Phillip Rogaway. Everything provable is provable in zero-knowledge. In Shafi Goldwasser, editor, *CRYPTO*, volume 403 of *Lecture Notes in Computer Science*, pages 37–56. Springer, 1988.
- [PV07] Rafael Pass and Muthuramakrishnan Venkatasubramanian. An efficient parallel repetition theorem for arthur-merlin games. In David S. Johnson and Uriel Feige, editors, *STOC*, pages 420–429. ACM, 2007.
- [Bar01] Boaz Barak. How to go beyond the black-box simulation barrier. In *FOCS*, pages 106–115, 2001.
- [PRS02] Manoj Prabhakaran, Alon Rosen, and Amit Sahai. Concurrent zero knowledge with logarithmic round-complexity. In *FOCS*, pages 366–375, 2002.
- [PTW11] Rafael Pass, Wei-Lung Dustin Tseng, and Douglas Wikström. On the composition of public-coin zero-knowledge protocols. *SIAM J. Comput.*, 40(6):1529–1553, 2011.
- [DGS09] Yi Deng, Vipul Goyal, and Amit Sahai. Resolving the simultaneous resettability conjecture and a new non-black-box simulation strategy. In *FOCS*, pages 251–260. IEEE Computer Society, 2009.

On the practical exploitability of Dual EC DRBG in TLS implementations

TANJA LANGE

(joint work with S. Checkowoy, R. Nicolerhagen, M. Fredrikson, A. Everspaugh, M. Green, T. Ristenport, J. Moskiewicz, and H. Shochom)

Dual EC DRBG is a pseudorandom number generator, which is included in standards NIST SP800-90, ANSI x9.82, and ISO 18031. Already during the public comment phase, concerns about its suitability were raised and one year later researchers presented a way that the designer could be hiding a backdoor in the system. Since the September-2013 publication of a document by Edward Snowden, it is clear that the NSA had designed the RNG with the backdoor in mind. This talk analyses how the backdoor can be exploited in practical applications, such as TLS on the internet. This involves interesting computations on elliptic curves.

More information on <http://projectbullrun.org/dual-ec>.

Fully Key-Homomorphic Encryption, Arithmetic Circuit ABE, and Compact Garbled Circuits

SERGEY GORBUNOV

(joint work with Dan Boneh, Craig Gentry, Shai Halevi, Valeria Nikolaenko, Gil Segev, Vinod Vaikuntanathan, Dhinakaran Vinayagamurthy)

(Key-policy) attribute-based encryption [1, 2] is a public-key encryption mechanism where every secret key sk_f is associated with some function $f : \mathcal{X} \rightarrow \mathcal{Y}$ and an encryption of a message μ is labeled with a public attribute vector $\vec{x} \in \mathcal{X}$. The encryption of μ can be decrypted using sk_f only if $f(\vec{x}) = 0 \in \mathcal{Y}$. Intuitively, the security requirement is collusion resistance: a coalition of users learns nothing about the plaintext message μ if none of their individual keys are authorized to decrypt the ciphertext.

Attribute-based encryption (ABE) is a powerful generalization of identity-based encryption and fuzzy IBE and is a special case of functional encryption. It is used as a building-block in applications that demand complex access control to encrypted data, in designing protocols for verifiably outsourcing computations, and for single-use functional encryption.

The past few years have seen much progress in constructing secure and efficient ABE schemes from different assumptions and for different settings. The first constructions [2, 6, 7, 8] apply to predicates computable by Boolean formulas which are a subclass of log-space computations. More recently, important progress has been made on constructions for the set of all polynomial-size circuits: Gorbunov, Vaikuntanathan, and Wee [9] gave a construction from the Learning With Errors (LWE) problem and Garg, Gentry, Halevi, Sahai, and Waters [10] gave a construction using multilinear maps. In both constructions the policy functions are represented as Boolean circuits composed of fan-in 2 gates and the secret key size is proportional to the *size* of the circuit.

We construct the first (key-policy) attribute-based encryption (ABE) system with short secret keys: the size of keys in our system depends only on the depth of the policy circuit, not its size. Our constructions extend naturally to arithmetic circuits with arbitrary fan-in gates thereby further reducing the circuit depth. Building on this ABE system we obtain the first reusable circuit garbling scheme that produces garbled circuits whose size is the same as the original circuit *plus* an additive $\text{poly}(\lambda, d)$ bits, where λ is the security parameter and d is the circuit depth. Save the additive $\text{poly}(\lambda, d)$ factor, this is the best one could hope for. All previous constructions incurred a *multiplicative* $\text{poly}(\lambda)$ blowup. As another application, we obtain (single key secure) functional encryption with short secret keys.

We construct our attribute-based system using a mechanism we call *fully key-homomorphic encryption* which is a public-key system that lets anyone translate a ciphertext encrypted under a public-key \vec{x} into a ciphertext encrypted under the public-key $(f(\vec{x}), f)$ of the same plaintext, for any efficiently computable f . We

show that this mechanism gives an ABE with short keys. Security is based on the subexponential hardness of the learning with errors problem.

We also present a second (key-policy) ABE, using multilinear maps, with short ciphertexts: an encryption to an attribute vector \vec{x} is the size of \vec{x} plus $\text{poly}(\lambda, d)$ additional bits. This gives a reusable circuit garbling scheme where the size of the garbled input is short, namely the same as that of the original input, *plus* a $\text{poly}(\lambda, d)$ factor.

It remains an interesting open problem to construct attribute-based encryption where the parameters do not depend on the depth of the circuit. This will lead to corresponding improvements in the parameters of the (reusable) garbling schemes. Also, we expect that fully-key homomorphic encryption can be used in interesting and versatile applications, and leave it as an open problem to find these applications.

REFERENCES

- [1] A. Sahai and B. Waters. Fuzzy identity-based encryption. In EUROCRYPT, 2005.
- [2] V. Goyal, O. Pandey, A. Sahai, and B. Waters. Attribute-based encryption for fine-grained access control of encrypted data. In ACM CCS, 2006.
- [3] A. Shamir. Identity-based cryptosystems and signature schemes. In CRYPTO, 1984.
- [4] C. Cocks. An identity based encryption scheme based on quadratic residues. In IMA Int. Conf., 2001.
- [5] D. Boneh and M. K. Franklin. Identity-based encryption from the Weil pairing. SIAM Journal on Computing, 2003. Preliminary version in CRYPTO '01.
- [6] A. B. Lewko, T. Okamoto, A. Sahai, K. Takashima, and B. Waters. Fully secure functional encryption: Attribute-based encryption and (hierarchical) inner product encryption. In EUROCRYPT 2010.
- [7] A. B. Lewko and B. Waters. New proof methods for attribute-based encryption: Achieving full security through selective techniques. In CRYPTO 2012.
- [8] X. Boyen. Attribute-based functional encryption on lattices. In TCC 2013.
- [9] S. Gorbunov, V. Vaikuntanathan, and H. Wee. Attribute-based encryption for circuits. In STOC, 2013.
- [10] S. Garg, C. Gentry, S. Halevi, A. Sahai, and B. Waters. Attribute-based encryption for circuits from multilinear maps. In CRYPTO, 2013.

Participants

Dr. Benny Applebaum

Department of Electrical
Engineering Systems
Tel Aviv University
Ramat Aviv 69978
ISRAEL

Aloni Cohen

MIT CSAIL
The Stata Center
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Prof. Dr. Daniel J. Bernstein

Department of Computer Science
University of Illinois at Chicago
M/C 249, 322 SEO
851 S. Morgan Street
Chicago IL 60607-7045
UNITED STATES

Dr. Özgür Dagdelen

Fachbereich Informatik
Technische Universität Darmstadt
64283 Darmstadt
GERMANY

Nina Bindel

Fachbereich Informatik
Technische Universität Darmstadt
64283 Darmstadt
GERMANY

Prof. Dr. Jintai Ding

University of Cincinnati
McMicken College of Arts & Sciences
7148 Edwards One
Cincinnati 45221-0037
UNITED STATES

Dr. Elette Boyle

Department of Computer Science
and Applied Mathematics, Technion
The Institute of Technology
Haifa 32000
ISRAEL

Dr. Cynthia Dwork

Microsoft Research
Silicon Valley
1065 La Avenida
Mountain View CA 94043
UNITED STATES

Prof. Dr. Johannes Buchmann

Fachbereich Informatik
Technische Universität Darmstadt
64283 Darmstadt
GERMANY

Rachid El Bansarkhani

Fachbereich Informatik
Technische Universität Darmstadt
64283 Darmstadt
GERMANY

Prof. Dr. Ran Canetti

Computer Science Department
School of Mathematics
Tel Aviv University
P.O. Box 39040
Tel Aviv 699 7801
ISRAEL

Prof. Dr. Marc Fischlin

Fachbereich Informatik
Technische Universität Darmstadt
64283 Darmstadt
GERMANY

Dr. Steven Galbraith

Department of Mathematics
The University of Auckland
Private Bag 92019
Auckland 1142
NEW ZEALAND

Prof. Dr. Shafi Goldwasser

MIT CSAIL
The Stata Center
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Florian Göpfert

Fachbereich Informatik
Technische Universität Darmstadt
64283 Darmstadt
GERMANY

Dr. Sergey Gorbunov

Department of Computer Science
MIT, Theory Division, Rm G32-578
32 Vassar Street
Cambridge, MA 02139
UNITED STATES

Dr. Vipul Goyal

Microsoft Research India
#9, Vigyan
Lavelle Road
Bangalore 560 001
INDIA

Dr. Iftach I. Haitner

School of Computer Science
Tel Aviv University
Schreiber Bldg., office 20
P.O. Box 39040
Ramat Aviv, Tel Aviv 699 78
ISRAEL

Prof. Dr. Dennis Hofheinz

Karlsruher Institut f. Technologie (KIT)
Arbeitsgruppe Kryptographie &
Sicherheit
Am Fasanengarten 5
76131 Karlsruhe
GERMANY

Justin Holmgren

MIT CSAIL
The Stata Center
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Prof. Dr. Yuval Ishai

Computer Science Department
TECHNION
Israel Institute of Technology
Haifa 32000
ISRAEL

Dr. Abhishek Jain

Department of Computer Science
MIT & BU - Theory Division
32 Vassar Street
Cambridge MA 02139
UNITED STATES

Prof. Dr. Antoine Joux

UPMC - Paris VI
LIP 6
4, Place Jussieu
75005 Paris Cedex
FRANCE

Dr. Yael Kalai

Microsoft Research
Office 14063
One Memorial Drive
Cambridge MA 02142
UNITED STATES

Prof. Dr. Tanja Lange

Dept. of Mathematics & Computer
Science
Eindhoven University of Technology
P.O. Box 513
5600 MB Eindhoven
NETHERLANDS

Prof. Dr. Hendrik W. Lenstra

Mathematisch Instituut
Universiteit Leiden
Postbus 9512
2300 RA Leiden
NETHERLANDS

Prof. Dr. Rachel Lin

Department of Computer Science
University of California
Santa Barbara, CA 93106
UNITED STATES

Prof. Dr. Vadim Lyubashevsky

École Normale Supérieure
Département d'informatique
45, rue d'Ulm
75230 Paris Cedex 05
FRANCE

Prof. Dr. Daniele Micciancio

Department of Computer Science &
Engineering
University of California, San Diego
Mail Code 0404
La Jolla CA 92093-5004
UNITED STATES

Prof. Dr. Jörn Müller-Quade

Karlsruher Institut f. Technologie (KIT)
Institut für Kryptographie und Sicherheit
Am Fasanengarten 5
76131 Karlsruhe
GERMANY

Prof. Dr. Moni Naor

Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O.Box 26
Rehovot 76100
ISRAEL

Prof. Dr. Rafail Ostrovsky

Henry Samueli School of
Engineering and Applied Science
Center for Information and Computation
Security
Los Angeles, CA 90095
UNITED STATES

Omer Paneth

Department of Computer Science
Boston University
808 Commonwealth Avenue
Boston MA 02215
UNITED STATES

Prof. Dr. Rafael Pass

Computer Science Department
Cornell University
White Hall
Ithaca, NY 14853
UNITED STATES

Dr. Chris Peikert

School of Computer Science
Georgia Institute of Technology
Atlanta, GA 30332-0160
UNITED STATES

Prof. Dr. Krzysztof Pietrzak

Institute of Science & Technology
(IST Austria)
Am Campus 1
3400 Klosterneuburg
AUSTRIA

Prof. Dr. Alon Rosen
The Interdisciplinary Center
Kanfei Nesharim St
P.O. Box 167
46150 Herzliya
ISRAEL

Dr. Guy Rothblum
Microsoft Research
One Microsoft Way
Redmond, WA 98052-6399
UNITED STATES

Ron Rothblum
Department of Computer Science
and Applied Mathematics
The Weizmann Institute of Science
P.O.Box 26
Rehovot 76100
ISRAEL

Prof. Dr. Claus-Peter Schnorr
Institut für Mathematik
J.W.Goethe-Universität
60054 Frankfurt am Main
GERMANY

Prof. Dr. Alice Silverberg
Department of Mathematics
University of California, Irvine
Irvine, CA 92697-3875
UNITED STATES

Prof. Dr. Nigel Smart
Department of Computer Science
University of Bristol
Merchant Venturers Bldg.
Woodland Road
Bristol BS8 1UB
UNITED KINGDOM

Prof. Dr. Stefano Tessaro
Department of Computer Science
University of California
Santa Barbara, CA 93106
UNITED STATES

Prof. Dr. Daniel Wichs
College of Computer Science
Northeastern University
215 Cullinane Hall
Boston, MA 02115
UNITED STATES

Mark Zhandry
Department of Computer Science
Stanford University
Stanford, CA 94305
UNITED STATES