THE FIELDS INSTITUTE
FOR RESEARCH IN MATHEMATICAL SCIENCES

ABSTRACTS 1.2

ADI AKAVIA
MIT

On Basing One-Way Functions on NP-Hardness

One-way functions are the cornerstone of modern cryptography. Informally speaking, one-way functions are functions that are easy to compute but are hard to invert (on the average case). There are several candidate functions, such as RSA or discrete-log, that are believed to be one-way, nonetheless, to date, no function was proved to be one-way. A puzzling question of fundamental nature is what are the minimal assumptions required for proving that a function is one-way. A necessary condition is that P does not equal NP (or more precisely, BPP does not equal NP, namely, that there is a problem in NP that cannot be solved by any probabilistic polynomial time algorithm). We ask whether this is also a sufficient condition. Namely, we ask whether there can be an efficient reduction from NP (that is, from the task of deciding an NP-complete language on the worst case) to a one-way function (that is, to the task of inverting a one-way function on the average case).

We proved two results on the impossibility of reducing NP to a one-way function; both results hold under the (widely believed) complexity assumption that coNP is not contained in AM. 1. There cannot be a reduction not even an adaptive reduction from NP to a ”size verifiable” one-way function; where we call f size-verifiable if, given y, the number of pre-image $|f^{-1}(y)|$ is efficiently computable, or, more generally, efficiently verifiable via an AM protocol. 2. There cannot be a non-adaptive reduction from NP to any one-way function (be it size-verifiable or not).

Our results improve on previously known negative results of [Feigenbaum-Fortnow,Bogdanov-Trevisan] by (i) handling adaptive reductions (whereas previous works were essentially confined to non-adaptive reductions), and by (ii) relying on a seemingly weaker complexity assumption.

In the course of proving the above results, we designed a new constant round interactive protocol for proving upper bounds on the sizes of NP sets. We believe this protocol may be of independent interest.

(Joint work with Oded Goldreich, Shafi Goldwasser and Dana Moshkovitz)

AMOS BEIMEL
Ben Gurion University and UCLA

Private approximation of search problems

Many approximation algorithms have been presented in the last decades for hard search problems. The focus of this paper is on cryptographic applications, where it is desired to design algorithms which do not leak unnecessary information. Specifically, we are interested in private approximation algorithms – efficient algorithms whose output does not leak information not implied by the optimal solutions to the search problems. Privacy requirements add constraints on the approximation algorithms; in particular, known approximation algorithms usually leak a lot of information.

For functions, [Feigenbaum et al., ICALP 2001] presented a natural requirement that a private algorithm should not leak information not implied by the original function. Generalizing this requirement to search problems is not straightforward as an input may have many different outputs. We present a new definition that captures a minimal privacy requirement from such algorithms – applied to an input instance, it should not leak any information that is not implied by its collection of exact solutions. Although our privacy requirement seems minimal, we show that for well studied problems, as vertex cover, max exact 3SAT, and clustering problems, private approximation algorithms are unlikely to exist even for poor approximation ratios. Similar to [Halevi et al., STOC 2001], we define a relaxed notion of approximation algorithms that leak (little) information, and demonstrate the applicability of this notion by showing near optimal approximation algorithms for max exact 3SAT that leak little information.

DANIEL J. BERNSTEIN
University of Illinois at Chicago

Proving tight security for Rabin-Williams signatures

Variants of the Rabin-Williams public-key signature system have, for twenty-five years, held the speed records for signature verification. Are these systems secure?

There are many other signature systems of RSA/Rabin type. One can break each system by factoring the signer's public key n or by breaking the system's hash function H. Are there other attacks? This is not an idle concern: some RSA-type systems have been broken by embarrassing attacks that (1) are much faster than known methods to factor n and (2) work for every function H (or a large fraction of choices of H), given oracle access to H.

Some systems have been proven immune to embarrassing attacks. For example, in the 1993 paper that popularized this line of work (along with the terminology "secure in the random-oracle model," which seems to have engendered considerable confusion), Bellare and Rogaway proved the following security property for the traditional "FDH" form of exponent-e RSA: every H-generic attack on RSA-FDH can be converted (without serious loss of efficiency) into an algorithm to compute eth roots modulo n.

Unfortunately, a closer look reveals that most of these proofs merely limit the embarrassment, without actually ruling it out. For example, the Bellare-Rogaway root-finding algorithm has only a $1/h$ chance of success, where h is the number of hash values seen by the FDH attack. Coron introduced a better algorithm having a $1/s$ chance of success, where s is the number of signatures seen by the FDH attack; but s can still be quite large.

Randomized signatures, in which long random strings are prepended to messages before the messages are signed, allow much tighter proofs. For example, every H-generic attack on randomized exponent-e RSA (or Rabin's 1979 signature system) can be converted into a algorithm to compute eth roots modulo n (or to factor n) with a good chance of success. But generating random strings takes time, and transmitting the strings consumes bandwidth. Can we do better?

A 2002 theorem of Coron is widely interpreted as saying that FDH is stuck at $1/s$: tight proofs require randomization. The randomized "RSA-PSS" and "PRab" systems have tight security proofs and work around the bandwidth problem, but they still take time to generate long random strings, and they are incompatible with the fastest known verification techniques. A tight security proof by Katz and Wang allows much shorter random strings for some RSA variants but breaks down for Rabin-Williams.

In my talk I'll explain three state-of-the-art variants of the Rabin-Williams public-key signature system. All three variants have tight security proofs, and all three provide extremely fast signature verification. What's most surprising is the FDH variant, which has a tight security proof despite hashing unrandomized messages; in the Rabin-Williams context, a minor technical assumption in Coron's theorem turns out to be a major loophole.

DAN BROWN
Certicom

Difficulty and Inapplicability of the RSA problem

Solving the RSA problem with a straight line program is shown to be almost as difficult as factoring. Proving - without random oracles - that breaking RSA-OAEP encryption is as difficult as the RSA problem, though, is shown to be infeasible.

DEBBIE COOK
Bell Labs/Lucent

Elastic block ciphers

We present a practical algorithm for creating variable-length block ciphers from existing block ciphers. The algorithm, actual constructions, performance and statistical analysis is discussed. A method for converting any key-recovery attack on the variable-length version of the cipher into an attack on the original cipher is presented. From this result, we conclude that the variable-length version is secure against key-recovery attacks if the original cipher is secure against such attacks. This is joint work with Moti Yung and Angelos Keromytis.

JEAN-MARC COUVEIGNES
Toulouse

Hard homogeneous spaces

Let G be a finite commutative group. A homogeneous space H for G is a finite set H of the same cardinality $S = \#H = \#G$ which is acted on simply transitively by G. This means that there is a single orbite and for any $g \in G$ not the identity, the permutation of H induced by g has no fixed points. If the left action is denoted by a dot we thus have

$$(\exists h \in H, g.h = h) \Longrightarrow g = 1.$$

We assume one can compute efficiently the composition law, inversion of an element and testing for equality. We also assume we can efficiently choose random elements in G and compute efficiently the action of G on H.

Now, because of of the lack of fixed points, there is a unique g mapping a given h_1 on a given h_2. We denote it by $\delta(h_2, h_1)$. We thus have

$$\delta(h_2, h_1).h_1 = h_2.$$

We may like to compute this $\delta(h_2, h_1)$. A related problem would be to complete a parallelogram namely, given $h_1, h_2, h_3 \in H$ compute the unique h_4 such that $\delta(h2, h1) = \delta(h4, h3)$.

We will be interested in Homogeneous Spaces for which these two problems are difficult. We call these Hard Homogeneous Spaces (HHS). Examples of HHS are provided by the discrete logarithm problem and many cryptographic protocols based on the discrete logarithm problem have a counterpart for any hard homogeneous space. And they are better discribed in this context.

I first discussed the notion of HHS in a 1997 talk and proposed new examples of HHS that do not come from DL problems. In this talk I will try to see if there is anything new to say about HHS almost ten years after.

GIOVANNI DI CRESCENZO
Telcordia

Virus Localization using Cryptographic Primitives

Virus detection is an important problem in the area of computer security. Modern techniques attempting to solve this problem fall into the general paradigms of signature detection and integrity checking. In this paper we focus on the latter principle, which proposes to label an executable or source file with a tag computed using a cryptographic hash function, which later allows the detection of any changes performed to the file. We suggest extending this principle so that not only changes to the file are detected, but also these changes are localized within the file; this is especially useful in the virus diagnostics which can then focus on the localized area in the file rather than the entire file. This implicitly defines an apparently new problem, which we call virus localization We design techniques to solve the virus localization problem based on repeated efficient applications of cryptographic hashing to carefully chosen subsets of the set of file blocks, for many of the most important and known virus infection techniques, which we characterize in our model.

JINTAI DING
Cincinnati and Darmstadt

Multivariate Quadratics for Hashing

We explore the idea of building a secure hash using quadratic or higher degree multivariate polynomials over a finite field as the Merkle-Damg aard compression function. We analyze the security properties and potential feasibility, there the compression functions are randomly chosen quadratic polynomials. Next, we propose to improve on the efficiency of the system by using some specially designed quadrctic polynomials with certain sparsity property, and the security of the system relies on stronger assumptions.

JUAN GARAY
Bell Labs

Searchable Symmetric Encryption: Improved Definitions and Efficient Constructions (This is joint work with Reza Curtmola, Seny Kamara, and Rafail Ostrovsky.)

Searchable symmetric encryption (SSE) allows a party to outsource the storage of his data to another party (a server) in a private manner, while maintaining the ability to selectively search over it. This problem has been the focus of active research for over a decade. In recent years, efficient solutions requiring only a constant number of rounds of interaction have been proposed at the expense of providing weaker security guarantees – specifically, these solutions allow the "access pattern" (i.e., the order in which the encrypted items are "touched" by the server) to be revealed. In this talk we consider this weaker security model as well, and show two constant-round solutions to SSE that simultaneously enjoy the following properties:

1) Both solutions are more efficient than previous constant-round schemes. In particular, the work performed by the server per returned document is constant as opposed to linear in the size of the data.

2) Both solutions enjoy stronger security guarantees than previous constant-round schemes. We point out shortcomings of previous notions of security for SSE, and show how to design constructions which avoid these pitfalls. Further, our second solution also achieves what we call *adaptive* SSE security, where queries into the database can be chosen adaptively (by the adversary) during the execution of the search.

We also consider *multi-user* SSE, where an arbitrary group of parties other than the owner can submit search queries. We formally define SSE in the multi-user setting, and present an efficient construction that achieves better performance than simply using access control mechanisms.

SHAI HALEVI
IRM Research

Mitigating Dictionary Attacks on Password-Protected Local Storage (Joint work with Ran Canetti, Shai Halevi and Michael Steiner)

We address the issue of encrypting data in local storage using a key that is derived from the user's password. The typical solution in use today is to derive the key from the password using a cryptographic hash function. This solution provides relatively weak protection, since an attacker that gets hold of the encrypted data can mount an off-line dictionary attack on the user's password, thereby recovering the key and decrypting the stored data.

we propose an approach for limiting off-line dictionary attacks in this setting without relying on secret storage or secure hardware. In our proposal, the process of deriving a key from the password requires the user to solve a puzzle that is presumed to be solvable only by humans (e.g, a CAPTCHA). We describe a simple protocol using this approach: many different puzzles are stored on the disk, the user's password is used to specify which of them need to be solved, and the encryption key is derived from the password and the solutions of the specified puzzles. Completely specifying and analyzing this simple protocol, however, raises a host of modeling and technical issues, such as new properties of human-solvable puzzles and some seemingly hard combinatorial problems. Here we analyze this protocol in some interesting special cases.

DAVID JAO
Waterloo

Isogenies as a cryptographic primitive

Isogenies between elliptic curves have long played a central role in the study of elliptic curve cryptography. Explicit calculations of isogenies have proved to be a crucial tool in elliptic curve point counting algorithms and in the theoretical analysis of the elliptic curve discrete logarithm problem. In this talk we explain how the efficient computation of isogenies provides a foundation for new implementations of cryptographic primitives sharing many of the traditional desirable properties of elliptic curve based cryptosystems. We also discuss recent results pertaining to the computation of isogenies and the implications of these developments in the context of constructing cryptographic primitives.

JONATHAN KATZ
Maryland

Concurrently-Secure Blind Signatures without Random Oracles or Setup Assumptions (Joint work with Camit Hazay, Chiu-Yuen Koo, and Yehuda Lindell)

We show a new protocol for blind signatures in which security is preserved even under arbitrarily-many concurrent executions. The protocol can be based on standard cryptographic assumptions, and is the first to be proven secure in a concurrent setting without random oracles or a common reference string.

AGGELOS KIAYIAS
Connecticut

Copyrighting public-key encryption and black-box traitor tracing

Copyrighting a function refers to the process of embedding hard-to-remove marks in the function's implementation while retaining its functionality. We present two methods for copyrighting discrete-log based public-key encryption functions and show how one can modularly obtain public-key traitor tracing schemes by composing copyrighted encryption with collusion secure code families. We provide a formalization of the notion of black-box traitor tracing that is the first that takes into account adversarially chosen plaintext distributions and we argue the security of our constructions in this setting. Regarding the modular approach, our constructions demonstrate how one can derive public-key traitor tracing by reducing the required "marking assumption" of collusion-secure codes to cryptographic hardness assumptions.

VLAD KOLESNIKOV
Bell-Labs/Lucent

How to tell which of the encrypted numbers is greater

We consider the problem of comparing two encrypted numbers and its extension – transferring one of the two secrets, depending on the result of comparison. Our constructions can be efficiently used in practice, e.g. in auctions with semi-honest auctioneer or purchasing digital goods, such as music or movies. We also define new primitives, which capture common security properties of one round protocols and computing on encrypted data in a variety of settings, which may be of independent interest.

HUGO KRAWCZYK
Techion and IBM Research

HMQV and Why Provable Security Matters

We outline the analysis of the HMQV protocol and use it as an example to show how much "provable security" matters to both theory and practice: Not only does provable security serve as a major tool for analyzing protocols under well-defined security properties but also serves as a design tool to achieve more secure, more efficient and simpler protocols. (And, as with any powerful weapon, much care needs to be exercised when using it...)

ANNA LYSYANSKAYA
Brown

Compact Ecash and Applications

The main idea of electronic cash is that, even though the same party (a Bank) is responsible for giving out electronic coins, and for later accepting them for deposit, the withdrawal and the spending protocols are designed in such a way that it is impossible to identify when a particular coin was spent. I.e., the withdrawal protocol does not reveal any information to the Bank that would later enable it to trace how a coin was spent. Since a coin is represented by data, and it is easy to duplicate data, an electronic cash scheme requires a mechanism that prevents a user from spending the same coin twice (double-spending), for example by identifying double-spenders and tracing all transactions that they have carried out.

In this talk, I will first present a scheme that allows a user to withdraw a wallet with W coins, such that the space required to store these coins, and the complexity of the withdrawal protocol, are proportional to logW, rather than to W. We achieve this without compromising the anonymity and unlinkability properties usually required of electronic cash schemes. We give a scheme that allows us to efficiently trace all coins that were spent by a double-spender. The security of our construction relies on a mix of cryptographic assumptions about groups with bilinear maps, and is in the random oracle model.

I will then show how to use the same methodology to achieve balance between accountability and privacy in other applications. In particular, we consider a setting where the amount of anonymous transactions with a particular merchant may be limited (e.g., so as to prevent money laundering). Finally, I will show that this methodology solves the problem of uncloneable group identification.

(Based on joint papers with Jan Camenisch, Susan Hohenberger, Markulf Kohlweiss, and Mira Meyerovich.)

ALEXANDER MAY
Paderborn

New Cryptanalytic Results for RSA with Small Exponents

The first part of the talk will provide a survey of attacks on RSA with special parameters. Restricted RSA parameters can be considered as a relaxation of the problems of taking e-th roots and the factorization problem, respectively. We give certain bounds under which the underlying problems become solvable in polynomial time. In the second part of the talk, we will present a new lattice-based attack on RSA with Small CRT-exponents. The attack is polynomial time whenever both d mod p-1 and d mod $q-1$ are smaller than $N^0.073$.

STEVEN MILLER
Rutgers

Provable Collisions in the Pollard Rho Algorithm for DLOG (Joint work with with David Jao and Ram Venkatesan)

The Pollard Rho algorithm for solving discrete logarithms works by iterating 3 types of group operations, finding a repeated value (collision), and then obtaining the desired exponent from the collision. It heuristically takes O(sqrt(p)) steps to find a collision, where p is the (prime) group order, and the collision is likewise expected to be nontrivial with probability 1-1/p. However, very little is known rigorously about the runtime of this algorithm. We prove a result using expander graphs that the collision time is $O(sqrt(p)(logp)^3)$ with probability arbitrarily close to one. The method also gives a result about the final step (of extracting the exponent from the collision) for almost all values of p. [Joint work with Ramarathnam Venkatesan, Microsoft Research]

KUMAR MURTY
University of Toronto

Factorization and Modular forms

We will discuss the possibility of using modular forms to factor integers. In particular, we will present a result indicating that the n-th Fourier coefficient of a Hecke eigenform tends to have factors in common with n with probability 1.

TATSUAKI OKAMOTO
NTT

An Efficient Public-Key Encryption under the Standard Assumptions

I will present a new method of constructing an efficient public-key and hybrid encryption schemes under the standard assumptions. Our schemes are almost as efficient as the Cramer-Shoup and Kurosawa-Desmedt schemes.

LEONID REYZIN
Boston University

Robust Fuzzy Extractors and Authenticated Key Agreement from Close Secrets (joint work with Yevgeniy Dodis, Jonathan Katz and Adam Smith)

Consider the problem of deriving high-quality keys from noisy data in the presense of an active adversary: two parties (or the same party at two different times), holding correlated secret inputs W and W', wish to agree on a uniformly distributed secret key R by sending a single message over an insecure channel. We study both the keyless case, where the parties share no additional secret information, and the keyed case, where the parties long-term secret SK that they can use to generate a sequence of session keys R_j using multiple pairs (W_j, W_j'). The former has applications to, e.g., biometric authentication, while the latter arises in, e.g., the bounded storage model with errors.

Our results improve upon previous work in several respects:

– The best previous solution for the keyless case with no errors (i.e., W=W') required the min-entropy of W to exceed 2n/3, where n is the bit-length of W. Our solution applies whenever min-entropy of W exceeds the minimal possible threshold n/2, and yields a longer key.

– Previous solutions for the keyless case in the presence of errors (i.e., W close, but not equal to, W') required random oracles. We give the first constructions in the standard model.

– Previous solutions for the keyed case were stateful, thus requiring synchronization between the parties. We give the first stateless solution.

REI SAFAVI-NAINI
Woollongong

Information Theoretic Bounds on Authentication Systems in Query Model (Joint Work with Peter Wild)

Authentication codes provide message integrity guarantees in an information theoretic sense within a symmetric key setting. Information theoretic bounds on the success probability of an adversary who has access to previously authenticated messages have been derived by Simmons and Rosenbaum, among others. We consider a strong attack scenario where the adversary is adaptive and has access to authentication and verification oracles. We derive information theoretic bounds on the success probability of the adversary and on the key size of the code. This brings the study of unconditionally secure authentication systems on a par with the study of computationally secure ones. We characterize the codes that meet these bounds and compare our result with the earlier ones.

BERRY SCHOENMAKERS
Eindhoven

Efficient Pseudorandom Generators Based on the DDH Assumption

A new family of pseudorandom generators based on the decisional Diffie-Hellman assumption is presented. The new construction is a modified and generalized version of the Dual Elliptic Curve generator proposed by Barker and Kelsey. Although the original Dual Elliptic Curve generator is shown to be insecure, the modified version is provably secure and very efficient in comparison with other pseudorandom generators based on discrete log assumptions. Our generator can be based on any group of prime order provided an additional requirement is met (i.e., there exists an efficiently computable function that in some sense ranks the elements of the group).

DOUGLAS R. STINSON
Waterloo

Hash Function Games and Two-Channel Authentication (Joint work with Atefeh Mashatan)

There has been some recent interest in authentication using two channels: an insecure broadband channel and an authenticated narrowband channel. This problem has been considered in the context of ad hoc networks, where there might not be a secret key shared by the two communicating parties, nor a public-key infrastructure to allow the use of public-key cryptography. For example, two devices might want to authenticate each other using an insecure wireless network, where the authenticated ”channel” consists of a human who manually transmits some small amount of authenticating data from one device to the other.

We study protocols for message authentication in this setting. We are interested in both interactive and noninteractive protocols. Our protocols are very simple and can be constructed from any suitable hash function. The security of our protocols depend on the difficulty of certain interactive hash function ”games”, the study of which seems to be an interesting problem in its own right.

RAM VANKATESAN
Microsoft Corporation

Randomized Sparse Representations of integer representations for Elliptic Curve and applications (Joint work with D.Jao, R.Raju)

We show that representing integers in base 0,1,x,y,..z with x,y,...,z chosen at random with mild conditions, leads to a sparse representation. We discuss the implications of this for performance and side channel attack models.

MOTI YUNG
Columbia University

Group Encryption

We present group encryption, a new cryptographic primitive which is the encryption analogue of a group signature. It possesses similar verifiability, security and privacy properties, but whereas a group signature is useful whenever we need to conceal the source (signer) within a group of legitimate users, a group encryption is useful whenever we need to conceal a recipient (decryptor) within a group of legitimate recievers. We introduce and model the new primitive and present sufficient as well as necessary conditions for its generic implementation. We then develop an efficient novel number theoretic construction for group encryption of discrete logarithms whose complexity is independent of the group size. To achieve this we construct a new public-key encryption for discrete logarithms that satisfies CCA2-key-privacy and CCA2-security. Applications of group encryption include settings where a user wishes to hide her preferred trusted third party or settings where verifiable well-formed ciphertexts are kept in a storage server that must be prevented from, both, learning the content of records and analyzing the identities of their retrievers.