Dagstuhl Seminar 02391:
CRYPTOGRAPHY
Sep. 22-27, 2002

Organizers:
Ueli Maurer[*],
Adi Shamir[†],
Jacques Stern[‡],
Moti Yung[§]

November 25, 2002

1 Summary

Since the advent of public key cryptography, about twenty-five years ago, the field of cryptography has been developing very rapidly. Constantly, there are new areas and new issues to investigate. The advance of the Internet as the major computing paradigm has increased the applicability and diversity of the field.

The aim of the 2002 Cryptography seminar was to provide an opportunity to focus on the scientific foundation of cryptography, to spot the emerging new areas based on recent advances in theory and technology needs, and to work on them.

The emphasis of the seminar was on the conceptual framework that allows the use of appropriate models, amenable to mathematical reasoning. Applications are natural in this field and were covered as well.

[*]ETH, Zurich, Switzerland
[†]Weizmann Inst., Rehovot, Israel
[‡]ENS, Paris, France
[§]Columbia U., New York, USA

We note that earlier cryptography seminars at Dagstuhl were held in the fall of 1992 and 1997. Similar workshops were also held at Luminy, France, and Monte-Verita, Switzerland. Previous meetings have led to valuable exchanges and to various investigations that advanced the field. The present seminar continued this tradition, with renewed topics, suitable to the current state of the art, and with concentration on a number of subjects that are being developed nowadays.

The following is a non-exhaustive list of topics discussed during the seminar:

1. Provable security of encryption and signature schemes.
2. New functions for cryptographic applications: the mathematics behind the Weil and Tate pairings over supersingular elliptic curves.
3. Novel cryptographic applications based on the bilinearity of pairings.
4. The applicability of various proof methodologies (random oracle proofs, generic model) to validation of cryptographic constructions.
5. New paradigms for cryptographic primitives (neural cryptography, quantum cryptography).
6. Methods for trust distribution.
7. Security models for multi-party protocols for function evaluation over private inputs and for commitment schemes: universal composability, non-malleability, etc.
8. New multi-party and commitment protocols.
9. Public key infrastructure and key distribution protocols.
10. Distributed cryptography over the Internet and in mobile networks (Byzantine agreement, threshold cryptography, fair exchange in ad-hoc mobile networks, etc.).
11. Relations between cryptographic primitives.
12. New methods in electronic voting.
13. Security in data retrieval.
14. New methods in content protection.
15. New notions in security: formal steganography, anonymity mechanisms.

16. New algebraic methods of cryptanalysis and their applications.

17. The effect of emerging developments in quantum computing on cryptographic primitives.

18. New design and analysis methodologies for, and experience with block ciphers (including the recent AES standard) and stream ciphers.

19. Improved efficiency of cryptographic mechanisms.

20. The influence of emerging computing environments and modern computer networks on cryptography.

21. The global implications of the "trusted computing platform" environments, recently proposed by the industry.

During the seminar, new directions for theoretical and applied research have been identified in numerous areas. The formal lectures, as well as the informal discussions, the moderated discussion session, and the informal session on recent results, were all inspiring and highly productive.

We feel that the subjects we have covered and worked on are most likely to influence cryptographic research in the coming few years. Furthermore, they have the potential to have impact on future applications of cryptographic techniques in computing systems.

2 Abstracts of Lectures

2.1 Blind Discrete Log Signatures and Uses of the Generic Model

Claus P.Schnorr

Consider the attacks on Schnorr blind signatures via a solution of the ROS-problem and Wagner's general birthday method. We show two limitations of this attack. Firstly, replacing the cyclic group G of prime order q by the s-fold direct product $G^{\times s}$ increases the work of generic interactive attacks against blind signatures to the s-power while increasing the work of the blind signature protocol merely by a factor s. Secondly, we support the conjecture that the best way to forge several additional signatures is to forge in separate attacks one additional signature per attack. We show that forging s additional signatures in a typical interleaved way requires work that grows as the s-power of the work for forging a single signature.

We present a generic variant GDSA$^+$ of the DSA in which the reduction function $f : G \to Z_q$ is replaced by a strong hash function. GDSA$^+$ is secure for a generic group G and a random hash function. We discuss a realistic use of the generic group model. Using the random encoding of group elements to further computations can be justified by applying in the real scheme a strong hash function to the group elements. We show that various recent conjectures can be proved in the generic group model.

2.2 Fractal Traversal of Merkle Trees

Markus Jakobsson

We introduce a technique for traversal of Merkle trees, and propose an efficient algorithm that generates a sequence of leaves along with their associated authentication paths. For one choice of parameters, and a total of N leaves, our technique requires a worst-case computational effort of 2logN hash function evaluations per output, and a total storage capacity of 8logN hash values.

This is joint work with Leighton, Micali and Szydlo.

2.3 Cryptography on a Quantum Computer

Claude Crépeau

The basic notion of Quantum Key Distribution will first be discussed. Then information theoretical notions of cryptography over quantum states such as encryption and authentication will be covered. In particular, we show that for quantum data, authentication imply encryption. Computational analogues will also be presented: quantum public-key cryptography, public-key authentication and impossibility of quantum digital signatures.

2.4 Relations Between the Complexity Assumptions Used in Tripartite Diffie-Hellman

Antoine Joux

Elliptic curves were first proposed as a tool for cryptography by V. Miller in 1985. Indeed, since elliptic curves have a group structure, they nicely fit as a replacement for more traditional groups in discrete logarithm based systems such as Diffie–Hellman or ElGamal. However, in 1993 Menezes, Okamoto and Vanstone showed that some special elliptic curves, called supersingular curves, are weaker than general elliptic curves. On these special curves, some additional properties allow an attacker to transport the discrete logarithm problem to a finite field where more efficient algorithms are available for discrete logarithm computation. Several recent papers have shown that the additional properties of weak curves can also be used positively. Indeed, it is possible to base cryptosystems on weak elliptic curves and to turn the additional properties of the curve into additional properties of the systems. Discussing the security hypothesis behind these systems is the main topic of this talk.

2.5 Framework for (Receipt-Free) Homomorphic Elections

Martin Hirt

Many homomorphic voting protocols in the literature are based on a concrete encryption scheme like ElGamal or Paillier, and a number of sub-protocols for the used encryption scheme are developed. In this talk, we present a framework for homomorphic elections: Given any semantically-secure homomorphic encryption scheme with some additional property (so called q-infertility), one can construct a very efficient secret-ballot voting scheme. All required sub-protocols can be developed generically for the encryption scheme. Furthermore, two specific receipt-free voting protocols for this framework are presented. The protocol with

randomizers is very efficient, but it achieves receipt-freeness only with respect to outside vote-buyers. The protocol based on ballot-shuffling requires a communication overhead which is linear in the number of authorities, but it guarantees receipt-freeness even with respect to a fixed minority of the authorities.

2.6 Distributing Trust on the Internet

Christian Cachin

This talk gives an overview of coordination protocols for asynchronous networks subject to Byzantine faults, such as binary and multi-valued Byzantine agreement, broadcast primitives like reliable and consistent broadcast, and atomic broadcast. Atomic broadcast immediately provides secure state-machine replication. The protocols are designed for an asynchronous wide-area network, such as the Internet, where messages may be delayed indefinitely, the servers do not have access to a common clock, and up to one third of the servers may fail in potentially malicious ways. Security is achieved through the use of threshold public-key cryptography, in particular through a cryptographic common coin based on the Diffie-Hellman problem.

2.7 Cryptanalysis of Stream Ciphers with Linear Masking

Shai Halevi

We describe a cryptanalytic technique for distinguishing some stream ciphers from a truly random process. Roughly, the ciphers to which this method applies consist of a "non-linear process" (say, akin to a round function in block ciphers), and a "linear process" such as an LFSR (or even fixed tables). The output of the cipher can be the linear sum of both processes. To attack such ciphers, we look for any property of the "non-linear process" that can be distinguished from random. In addition, we look for a linear combination of the linear process that vanishes. We then consider the same linear combination applied to the cipher's output, and try to find traces of the distinguishing property.

In this work we analyze two specific distinguishing properties. One is a linear approximation of the non-linear process, which we demonstrate on the stream cipher SNOW. This attack needs roughly 2^{95} words of output, with work-load of about 2^{100}. The other is a "low-diffusion" attack, that we apply to the cipher Scream-0. The latter attack needs only about 2^{43} bytes of output, using roughly 2^{50} space and 2^{80} time.

This is joint work with Don Coppersmith and Charanjit Jutla.

2.8 Efficient Computation Modulo a Shared Secret with Applications to The generation of Shared Safe-Prime Product

Jan Camenisch

We present a new protocol for efficient distributed computation modulo a shared secret. We further present a protocol to distributively generate a random shared prime or safe prime that is much more efficient than previously known methods. This allows one to distributively compute shared RSA keys, where the modulus is the product of two safe primes, much more efficiently than was previously known.

2.9 A New Knapsack Cryptosystem

Tatsuaki Okamoto

Although extensive research has been made by numerous cryptographers and mathematicians to realize the concept of public-key cryptosystems for more than 20 years, very few concrete techniques that seem to be secure have been found.

A typical approach is based on number theoretic tricks such as RSA-Rabin, Diffie-Hellman-ElGamal, and the elliptic curve cryptosystems. Currently this approach is the most successful and these schemes are employed in many applications.

Another typical approach is based on combinatorics tricks such as knapsack cryptosystems and lattice based systems. Unfortunately, this approach is not considered to be successful, since almost all schemes based on this approach have been broken.

It is, however, very important to realize a new practical public-key cryptosystem based on combinatorics tricks from the following reasons (in the lights of security and practice):

- (Security) Number theoretic problems might be broken in the future drastically (e.g., by new algorithms or new machines such as quantum computers). Combinatorics tricks may provide us a more secure trapdoor trick (which may be almost as secure as the NP hard problems).
- (Practice) Combinatorics tricks may provide us a much faster (efficient) cryptosystem than number theoretic cryptosystems.

This talk presents a new promising approach of realizing a public-key cryptosystem based on a knapsack (subset-sum) problem. This approach employs the ring of integers, O_K, of an algebraic number field, K, which provides us a huge number of freedom of choosing trapdoor parameters.

This is joint work with Keisuke Tanaka and Shigenori Uchiyama.

2.10 Random Oracles May Be Subtle: The Density of e-th Residues

Jacques Stern

ESIGN is an efficient signature scheme that has been proposed in the early nineties. Recently, an effort was made to lay ESIGN on firm foundations, using the methodology of provable security. A security proof in the random oracle model appeared in support for ESIGN. However, several unexpected difficulties were found. Firstly, it was observed that the proof holds in a more restricted model of security than claimed. Even if it is quite easy to restore the usual security level, this shows that the methodology of security proofs is more subtle than it at first appears. Secondly, it was found that the proof needs the additional assumption that e is prime to $\varphi(n)$, thus excluding the case where e is a small power of two, a very attractive parameter choice. The difficulty here lies in the simulation of the random oracle, since it relies on the distribution of e-th powers, which is not completely understood from a mathematical point of view, at least when e is not prime to $\varphi(n)$. In this talk, we prove that the set of e-th power modulo an RSA modulus n, which is a product of two equal size integers p,q, is almost uniformly distributed on any large enough interval. This property allows to complete the security proof of ESIGN. We actually offer two proofs of our result: one is based on two-dimensional lattice reduction, and the the other uses Dirichlet characters. Besides yielding better bounds, the latter is one of the few examples of the use of analytic number theory in cryptography.

This is joint work with T. Okamoto.

2.11 Universally Composable Protocols: Some recent results

Ran Canetti

Recently, a new paradigm for defining security of cryptographic protocols was proposed. The salient property of notions of security that follow this paradigm

(called universally composable (UC) security) is that they guarantee security even when a protocol is used as a component within an arbitrary system. In particular, UC notions guarantee security even when an unbounded number of protocol instances are running concurrently in an adversarially controlled manner, they guarantee strong non-malleability with respect to arbitrary protocols, and more. Such properties are crucial for arguing the security of cryptographic protocols in complex and unpredictable environments such as the Internet.

The first part of the talk reviews the general methodology for formulating UC notions of security, and the composition theorem that underlies the strong security properties provided. We then quickly survey a number of works that use the UC paradigm in a variety of settings and purposes. The second part describes in more detail three of these works. The first (joint with Lindell, Ostrovsky and Sahai) demonstrates how to realize any multiparty functionality with any number of faults, in the common reference string model. The second (joint with Krawczyk) provides UC notions of secure key-exchange and secure channels and shows how to realize them. The third work (joint with T. Rabin) provides a new tool that plays a crucial role in the two previous works. This tool, called universal composition with joint state, enables us to deal with protocols where multiple instances use some joint module but are otherwise disjoint. It may well be useful elsewhere.

2.12 Neural Cryptography

Adi Shamir

In this talk we describe and analyze a new key exchange protocol proposed by Kanter Kinzel and Kanter, which is based on mutually learning neural networks. This is a new potential source for public key cryptographic schemes which are not based on number theoretic functions, and have small time and memory complexities. In the first part of the talk we analyze the properties of the scheme, explain why the two parties converge to a common key, and why an attacker using a similar neural network is unlikely to converge to the same key. However, in the second part of the talk we show that this key exchange protocol can be broken in three different ways, and thus it is completely insecure.

This is joint work with Alexander Klimov and Anton Mityagin.

2.13 Cryptography with Guardia Angels: Bringing Civilization to Pirates

Serge Vaudenay

In contrast with traditional cryptographic protocols in which parties can have access to common third parties, and where at least one of them is assumed to be honest, we propose here a new model which is relevant for networks of communication devices with security modules. We then focus on the problem of fair exchange in this model. We propose a probabilistic protocol which provides arbitrarily low unfairness (involving a complexity cost).

This is joint work with Gildas Avoine.

2.14 Efficient Non-Interactive, Reusable and Non-Malleable Commitments

Ivan Damgård

Non-malleable commitments secure according to the definition used in recent literature are not necessarily secure if the adversary is allowed to see more than one commitment from honest players before having to make his own, or if he is allowed to produce as output more than one commitment. We give concrete examples showing that both the multiple inputs and multiple outputs case lead to strictly stronger notions, both for unconditionally hiding and for unconditionally binding commitments. We point out that the security proofs for known non-interactive unconditionally hiding schemes fail in the multiple input case. We then present new schemes, based on the discrete logarithm or the strong RSA assumptions, that are efficient, non-interactive and unconditionally hiding, and which remain secure even if several commitments are known to the adversary, and if he may produce an arbitrary number of output commitments. We also show an application of our scheme to improve the universally composable scheme of Damgård and Nielsen, so that it can work with a reference string of size independent of the number of players.

2.15 Breaking the $O(n^{1/(2k-1)})$ Barrier for Information-Theoretic Private Information Retrieval

Yuval Ishai

Private Information Retrieval (PIR) protocols allow a user to retrieve a data item from a database while hiding the identity of the item being retrieved. PIR was studied in two main settings: the information-theoretic and the computational setting. In an *information-theoretic, k-server PIR protocol* the database is

replicated among k servers, and each server learns nothing about the item the user retrieves. The cost of such protocols is measured by the *communication complexity* of retrieving one out of n bits of data. For any fixed k, the best protocols prior to our work had complexity $O(n^{1/(2k-1)})$ (Ambainis, 1997). Since then several methods were developed in an attempt to beat this bound, but all these methods obtained the same asymptotic bound. In this work, this barrier is finally broken and the complexity of information-theoretic k-server PIR is improved to $n^{O((log \log k)/(klogk)}$. The new PIR protocols can also be used to construct k-query *locally decodable codes* with length $\exp[n^{O((\log \log k)/(klogk))}]$ over a binary alphabet, compared to $\exp[n^{(1)/(k-1)}]$ in previously known constructions. The improvements presented in this paper apply even for small values of k: the PIR protocols are more efficient than previous ones for every $k \geq 3$, and the locally decodable codes are shorter for every $k \geq 4$.

This is joint work with Amos Beimel, Eyal Kushilevitz, and Jean-Francois Raymond.

2.16 Some Recent Results and Open Problems for Single Database PIR

Rafail Ostrovsky

In this talk I described several recent results relating to Single Database Private Information Retrieval (PIR) protocol. In particular, I discussed (1) the complexity of PIR protocol based on trapdoor one-way permutations and some open problems and (2) applicability of a new a notion of a "batch code" in order to achieve amortized savings (with 0-error) in the database work.

Part of the talk is based on a joint work with Yuval Ishai, Eyal Kushilevitz and Amit Sahai.

2.17 Secret-Ballot Receipts and Transparent Integrity

David Chaum

Receipts showing exactly who you voted for – just what is generally wanted and expected today – have been outlawed to prevent vote selling and other abuses. A new kind of receipt cannot be abused. It also lets you be sure that your votes are correctly included in the final tally, even if all the computers used to run the election are compromised! Receipts are printed on two-layer media by a modified version of familiar receipt printers. You can read them clearly in the booth; but before leaving, you must separate the layers and choose which one to keep.

Either one you take has coded in it the vote information you saw, though your choices can now only be read using keys divided among computers run by election officials. The layer you take is supplied by the voting machine for publication on an official election website, where you can verify that it is posted. After deriving the tally from the posted receipts, a lotto-like draw selects parts that must be decrypted for inspection, but not so many parts that privacy is compromised. Anyone with a computer can simply check all the decryptions, which should also be published on the website, and thereby verify that the final tally must be correct. The printers and media are practical and under development. The overall system cost is lower than with today's voting machines and the hardware can additionally be used for other purposes year round. Current election system functionality, including write-ins and provisional ballots, is fully supported and can be extended significantly. A variety of public policy issues are raised. (See www.vreceipt.com.)

2.18 Designing Group Signatures from Traitor Tracing Schemes

Moti Yung

Digital Signature emerges naturally from Public-Key Encryption based on trapdoor permutation, and the "duality" of the two primitives was noted as early as Diffie-Hellman's seminal work. In this work we make the point that two well known cryptographic primitives whose connection has not been noticed so far in the literature enjoy an analogous duality. The primitives are Group Signature Schemes (used for anonymity of signers) and (Public-Key) Traitor Tracing (used in the DRM area). Based on the observed duality, we solve a number of open problems regarding the design of Group Signatures, the first such schemes based on the DDH problem where the signature size is sub-linear in the group size; the design is based on translating a traitor tracing scheme based on the discrete logarithm representation into a group signature. We also show how traceability codes (a natural notion in traitor tracing scheme) leads to the design of certain group signature schemes (with short-signature) based on the existing group signatures where signature is of linear size.

This is joint work with Aggelos Kiayias.

2.19 Tight Bounds for Shared Memory Systems Accessed by Byzantine Processes

Rebecca N. Wright

We provide efficient constructions and strong lower bounds for shared memory systems accessed by n processes, up to t of which may exhibit Byzantine faults, in a model previously explored by Malkhi et al. (DISC 2000). We show that sticky bits are universal in the Byzantine failure model provided n is at least 3t+1, an improvement over the previous result requiring n to be at least (2t+1)(t+1). Our result follows from a new strong consensus construction that uses sticky bits and tolerates t Byzantine failures among n processes for any $n > 3t$. (This is the best possible bound on n for strong consensus.) We also present tight bounds on the efficiency of implementations of strong consensus objects from sticky bits and similar primitive objects.

This is joint work with Michael Merritt, Omer Reingold, and Gadi Taubenfeld, and will be appearing in DISC '02.

2.20 Robust Cryptography

Ross Anderson

In this talk I explore attacks on hardware implementations of cryptography, such as smartcards and cryptoprocessors, and defenses against them. We have shown that optical probing with laser or simple flash-lamps can induce faults that leak information; there is significant scope for developing more powerful variants using X-rays. We have also developed a prototype smartcard chip that stops such attacks and also attacks based on power analysis, by using redundant self-checking logic. I report the results of the first few months testing of this chip.

2.21 Trusted Computing

Ross Anderson

I led a discussion on the potential, and the problems, of the "trusted computing" initiative of Intel, Microsoft and others. It has the potential to improve security of systems in some respects, but poses problems for public policy, privacy and research.

2.22 Efficient Adaptively Secure Multiparty Computation from Threshold Homomorphic Encryption

Jesper Buss Nielsen

We show how to build general multiparty computation with adaptive security from homomorphic threshold cryptoystems with specific properties. One example of such a system is Paillier's cryptosystem. The protocol is essentially as efficient as a previous protocol by Cramer, Damgård and Nielsen that was only statically secure, and hence much faster than solutions using non-committing encryption. The protocol can be proved secure in Canetti's universally composable framework.

This is joint work with Ivan Damgård.

2.23 On Deniability in Quantum Key Exchange

Don Beaver

We show that claims of "perfect security" for keys produced by quantum key exchange (QKE) are limited to "privacy" and "integrity." Unlike a one-time pad, QKE does not necessarily enable Sender and Receiver to pretend later to have established a different key. This result is puzzling in light of Mayers' "No-Go" theorem showing the impossibility of quantum bit commitment. But even though a simple and intuitive application of Mayers' protocol transformation appears sufficient to provide deniability (else QBC would be possible), we show several reasons why such conclusions are ill-founded. Mayers' transformation arguments, while sound for QBC, are insufficient to establish deniability in QKE.

Having shed light on several unadvertised pitfalls, we then provide a candidate deniable QKE protocol. This itself indicates further shortfalls in current proof techniques, including reductions that preserve privacy but fail to preserve deniability. In sum, purchasing undeniability with an off-the-shelf QKE protocol is significantly more expensive and dangerous than the mere optic fiber for which "perfect security" is advertised.

2.24 Some Recent Work Constructing Block-Cipher Modes of Operation

Phillip Rogaway

In this talk I survey my recent work on the construction of block-cipher modes of operation. I look at the the schemes OCB [RBBK02], $\dot{O}CB$ [Ro02a], and

EMD [Ro02b]. For each mode I explain exactly what is the defined goal and why a scheme has been constructed for that goal. Namely, OCB is a nonce-based authenticated-encryption scheme; $\dot{O}CB$ is a nonce-based authenticated-encryption scheme allowing associated-data (i.e., an "authenticated header"); and EMD is a tweaked, strong pseudorandom permutation with a long blocksize (intended for disk-sector encryption). Each mode is aggressively optimized yet simple and provably secure.

2.25 Orders, Discrete Logarithms and Group Structures

Johannes Buchmann

It is known that any finite Abelian group G is a direct product of cyclic subgroups $H_1, \ldots, H_k$ such that the order of H_i divides the order of $H_{i+1}, 1 \leq i < k$. The number and the orders of the subgroups H_i are uniquely determined by G. In this talk we present an algorithm that computes such subgroups H_i and their orders using $O(\sqrt{G})$ group operations (multiplications, divisions, comparisons) and look ups containing $O(\sqrt{G})$ group elements.

2.26 Equivalence Between The Random Oracle Model and The Random Cipher Model

Jean-Sébastien Coron

We show that the random model is equivalent to the random cipher model. We show that any scheme secure in the random oracle model is also secure in the random oracle model, and conversely. We introduce a new notion of indistinguishability between constructions involving random oracles or random ciphers. This is of practical interest since this enables to build secure block ciphers with larger input/output blocks, as it is often required in some public-key schemes.

This is joint work with A. Joux and D. Pointcheval.

2.27 Efficient Construction of (Distributed) Verifiable Random Functions

Yevgeniy Dodis

We give the first simple and efficient construction of verifiable random functions (VRFs). VRFs, introduced by Micali et al. [MRV99], combine the proper-

ties of regular pseudorandom functions (PRFs) [GGM86] (i.e., indistinguishability from a random function) and digital signatures [GMR88] (i.e., one can provide an unforgeable proof that the VRF value is correctly computed). The efficiency of our VRF construction is only slightly worse than that of a regular PRF construction of Naor and Reingold [NR97]. In contrast to ours, the previous VRF constructions [MRV99,Lys02] all involved an expensive generic transformation from verifiable unpredictable functions (VUFs), while our construction is simple and direct.

We also provide the first construction of distributed VRFs. Our construction is more efficient than the only known construction of distributed (non-verifiable) PRFs [Nie02], but has more applications than the latter. For example, it can be used to distributively implement the random oracle model in a publicly verifiable manner, which by itself has many applications (e.g., constructing threshold signature schemes).

Our main construction is based on a new variant of decisional Diffie-Hellman (DDH) assumption on certain groups where the regular DDH assumption does not hold. We do not make any claims about the validity of our assumption (which we call sum-free DDH, or sf-DDH). However, this assumption seems to be plausible based on our current understanding of certain candidate elliptic and hyper-elliptic groups which were recently proposed for use in cryptography [JN01,Jou00]. We hope that the demonstrated power of our sf-DDH assumption will serve as a motivation for its closer study.

2.28 In How Many Ways Can You Write Rijndael?

Eli Biham

We ask the question what happens if we replace all the constants in the Advanced Encryption Standard (aka AES, Rijndael), including the replacement of the irreducible polynomial, the coefficients in the Mix Column operation, the affine transformation in the S box, etc. We showed that such replacements can create new dual ciphers – equivalent to the original in all aspects. We presented several such dual ciphers, such as the square of Rijndael, and dual ciphers with the irreducible polynomial replaced by a primitive polynomial. We also described another family of dual ciphers to which we call the logarithm of Rijndael, and discussed self-dual ciphers, and application of dual ciphers.

In addition we showed that a recent observation of Fuller and Millan that all the 8 output bits of the S box of Rijndael are represented only as one function $f : \{0,1\}^8 \rightarrow 0,1$ up to affine transformations is also applicable to many other ciphers, and that many ciphers not only have this property, but also use the same

bit-functions as other ciphers, up to affine transformations. We summarized by showing that the least significant bit of the output of S9 of Misty1 is not affected by rotations of the input, and that 5 bits of the 9 output bits are not affected when the input is rotated by one bit to the right and the output by one bit to the left.

This is joint work with Elad Barkan.

2.29 A Signature Scheme with Efficient Protocols

Anna Lysyanskaya

Digital signature schemes constitute a fundamental cryptographic primitive, of use both in its own right, and as a building block in cryptographic protocol design.

Let us look at the main two signature algorithms – the algorithm for signing, and the algorithm for verifying a signature – as protocols. Given a signature scheme, the signing protocol is between a user who wants his message m signed, and the signer with public key PK. The verification protocol is between a user who holds a signature $\sigma_P K(m)$ and a verifier who wants to make sure that the signature held by the user is valid.

Let us make these two protocols as secure as possible. That is to say, they must conform to their specification, and should not leak any more information than the specification allows. We want to make possible a signing protocol in which the signer does not learn anything about the message he is signing, and a verification protocol in which the verifier does not learn anything about the signature he is verifying, other than that it is a valid signature.

We propose a practical and provably secure signature scheme and show such secure protocols for signing and verifying a signature. The proposed signature scheme and protocols are secure under the strong RSA assumption, and are efficient enough to be useful in practice, as a building block for the design of anonymity-enhancing cryptographic systems, such as electronic cash, group signatures, and anonymous credential systems.

This is joint work with Jan Camenisch.

2.30 Reducing the Trivial Lattice

Mike Szydlo

We discuss a more narrow class of lattices than are typically considered in cryptography: lattices which are isomorphic to the 'Trivial Lattice': Z^n. If such

a lattice is presented in terms of the standard basis, its short vectors are evident. However, this need not be the case: an instance of the trivial lattice may also be defined via a Gram matrix of some integral basis. Such lattices, also called orthogonal, arise in the cryptanalysis of GGH and related signature schemes. In such schemes, we review how averaging a transcript of size 10,000 signatures, reduces the problem of private key recovery to that of the shortest vector problem in an orthogonal lattice. This work presents a new polynomial time oracle-algorithm to solve the shortest vector problem in orthogonal lattices. This oracle is required to solve (certain special cases of) the 'Lattice Distinguishing Problem', which we define. In light of this polynomial reduction, we discuss the feasibility of such an oracle. We continue with a construction to realize the oracle using the statistics encoded in lattice theta functions. Concretely, one may completely distribute the collection of a large number of medium length lattice vectors, and use them to decide whether 2 lattices are isomorphic. Completely different in approach from standard lattice reduction algorithms, these algorithms exploit integrality, and statistical properties of lattices. The results may also be interpreted as a description of the private information leakage present in signatures produced with the GGH paradigm.

2.31 Efficient Proven Secure Steganography

Yvo Desmedt

At Crypto 2002 Hopper-Langford-von Ahn presented several steganographic schemes. We:

1. explain why the security proof of their first scheme is incorrect and how it can be fixed.
2. demonstrate that their second scheme does not work for several distribution (where our schemes do work).
3. present several new schemes with 0% error. These schemes can be used when one desires:
   (a) perfect hiding, perfect privacy, perfect reliability, 0% error provided the probability distribution is known. In this case we can prove that our bandwidth is optimal. We prove necessary and sufficient conditions on the probability.
   (b) high bandwidth perfect hiding, perfect privacy, 0% error (but no perfect reliability). We assume the probability distribution is known.

(c) high bandwidth perfect hiding, perfect privacy, 0% error under a variant of the Hopper-Langford-von Ahn passive model. We only assume in this model that the sender or either the receiver (but not necessarily both as the CMU model assumes) have access to the random (not necessarily uniform) oracle.

(d) proven secure high bandwidth 0% error schemes using the Hopper-Langford-von Ahn passive model. This means sender and receiver do not know the distribution, but have access to a random (not necessarily uniform) oracle.

4. introduce a chosen cover distribution attack model

5. introduce a stego-key model, in which the key used by the sender and receiver must be hidden (chosen by a known distribution not necessarily uniform). We present a scheme to send binary messages.

2.32 Linear Secret Sharing and Efficient Reconstruction of the Secret from Corrupted Shares

Ronald Cramer

We show that strong multiplication, a certain non-linear property that a linear secret sharing scheme may enjoy, facilitates efficient reconstruction of a shared secret even if the shares contributed by some non-qualified set are corrupted. This is done by looking in the right way at a classical error-correction algorithm for Reed-Solomon codes, and observing that strong multiplication enables its generalization to linear secret sharing schemes.

We note that, as opposed to the case of linear secret sharing schemes with ordinary multiplication, it is an open problem in the algebraic theory of secure multi-party computation whether the complexity of so-called Q3 access structures when measured in terms of linear secret sharing schemes with strong multiplication realizing them, is polynomial in their complexity in terms of linear secret sharing schemes as such. This question is directly related to the existence of certain error-free secure distributed multiplication protocols.

2.33 Certificate-Based Encryption

Craig Gentry

We introduce the concept of certificate-based encryption, in which a certificate – or, more generally, a signature – acts not only as a certificate but also as a decryption key. To decrypt a message, a keyholder needs both his secret key and an up-to-date certificate from his CA (or a signature from an authorizer). Certificate-based encryption combines the best aspects of identity-based encryption (implicit certification) and public-key encryption (no escrow).

We demonstrate how certificate-based encryption can be used to construct a PKI requiring less infrastructure than previous proposals, including Micali's Novomodo, Naor-Nissim and Aiello-Lodha-Ostrovsky. In each of these previous proposals, an encrypter must check the message recipient's certification status before encrypting to him, and considerable infrastructure is needed to process these queries. With certificate-based encryption, third-party certificate queries are unnecessary since certification is implicit – if the message recipient is not certified, he will not be able to decrypt. Eliminating third-party queries has other benefits as well. In particular, it reduces the system's vulnerability to denial-of-service attacks, and it allows a more sensible business model (where a CA charges clients for certificates, and does not respond to queries from non-clients).

The PKI scheme makes use of Gentry and Silverberg's hierarchical identity-based encryption construction, and has transmission costs similar to Aiello-Lodha-Ostrovsky (if there were no third-party queries). Computational costs are higher (since the above proposals are all hash-based), but are still reasonable. Other applications and extensions are possible. For example, a certificate-based encryption scheme with very high time-granularity is feasible using a reversal of Katz's forward-secure encryption scheme. Also, a keyholder's ability to decrypt may be made contingent on its possession of multiple signatures from different signers on different documents without increasing decryption time or the length of the ciphertext (though encryption time increases).

2.34 A Computationally Secure Key Pre-Distribution System Tolerating Arbitrary Coalitions

Berry Schoenmakers

A key pre-distribution system (KDS) allows any pair of users to establish a shared secret key, which is unique to the pair of users and unknown to all other users. The private and public keys for each user are generated and distributed by a trusted party, and these keys all depend on a master key held by the trusted party. We present a KDS for which any coalition of users, pooling their private keys, is unable to find the key shared by pairs of users outside the coalition.

Hence, there is no upper bound on the size of coalitions as in the "classical" Blom's KDS. Our construction relies on the Weil pairing (or any of its relatives) for the computation of the shared keys. The secrecy of the shared keys relies on a computational assumption related to the Diffie-Hellman assumption–unlike Blom's scheme, which achieves unconditional secrecy.

This is joint work with Pim Tuyls.

2.35 Lifting Part of the Veil on the Murder Of Patrice Lumumba; or Breaking 1961 Hagelin Ciphertexts

Bart Preneel

In this talk we discuss some cryptanalytic work carried out for the Belgian Parlement in the Fall of 2001. The motivation for this work was the investigation carried out by the Parliament on the circumstances of the murder on Patrice Lumumba. We were provided with 15 enciphered telexes sent between December 1960 and February 1961 between Brussels and Elisabethville and Brazzaville. In addition, 5 likely plaintexts (with errors) were available. We describe how we were able to identify that the PRINTEX variant used was the Hagelin C-38. We also succeeded in recovering the key settings by improving the Morris algorithm published in 1978. We also identified and cryptanalyzed the mechanism to encrypt session keys (Playfair). One of the telexes, dating from a few days before the murder, revealed some interesting new information.

2.36 Multi-variate Public Key Schemes

Hans Dobbertin

We consider the HFE (Hidden Field Equations) scheme for general power functions. Two selected cryptanalytic results are presented: The system can be broken if (1) the exponent is a geometric series and if (2) the difference of two digits in the representation of the exponent is half of the extension degree of the considered field extension. In the first case we can compute a substitute for the private key, whereas in the second case we can even recover the private key.