DIMACS Technical Report 97-80
January 1998

DIMACS Research and Education Institute (DREI '97)
Cryptography and Network Security
July 28 - August 15, 1997
Abstracts of Talks Presented

by

Joan Feigenbaum
Research Program Director
AT&T Labs - Research

DIMACS is a partnership of Rutgers University, Princeton University, AT&T Labs-Research, Bell Labs and Bellcore.

DIMACS is an NSF Science and Technology Center, funded under contract STC–91–19999; and also receives support from the New Jersey Commission on Science and Technology.

Week 1 (July 27 - August 2, 1997): Underlying Mathematics

A High Speed, Software-Driven, Stream Cipher

Bill Aiello, Bellcore
aiello@bellcore.com

In this talk we describe a pseudo-random generator which is very fast in software. This generator can be used as a pseudo-random one-time pad to implement a high-speed, software driven, stream cipher. The generator has three parameters which control a speed vs space vs security tradeoff.

Our generator can be based on any one-way permutation, or block cipher primitive. All the known primitives which are believed to be secure are not fast enough in software for high-speed applications. The efficiency of the generator is achieved by mostly limiting the slow primitive operations to a preprocessing step which computes a cryptographically strong, pseudo-random table. The on-line computations consists mostly of a few simple operations such as table look-ups and additions. For example, the speeds attained for memory-to-memory encryption (i.e., assuming disk I/O is charged to the calling application) on a Pentium Pro with reasonable settings of the parameters is 70 megabits/sec.

Based only on the security of the underlying primitive, the generator achieves many security and statistical properties which we will discuss.

This is joint work with S. Rajagopalan (Bellcore) and R. Venkatesan (Microsoft).

Linear Algebra Mod N

Eric Bach, University of Wisconsin
bach@cs.wisc.edu, http://www.cs.wisc.edu/~bach/bach.html

Many computations in number theory rely on the solution of linear equations. An example is the final stage of a difference-of-squares integer factoring method such as the quadratic sieve, which requires one to solve a large linear system mod 2. In other cases, such as computing discrete logarithms, systems must be solved modulo composite numbers. Unfortunately, the lion's share of the literature on computational linear algebra assumes the coefficients come from a field, and gives short shrift to systems over finite rings such as Z/N. We will survey this latter area.

Efficient Methods for Modular Arithmetic

Josh Benaloh, Microsoft
benaloh@microsoft.com

As almost any student who has implemented large integer "BigNum" arithmetic routines will attest to, the hardest operation by far is the integer division with remainder operation required to do modular arithmetic. This talk will describe some common (and some not so common) methods for doing modular arithmetic as efficiently as possible. Particular emphasis will be given to efficient methods of modular reduction including the Montgomery method to bypass the division step entirely.

Cryptanalysis of the Portz Interconnection-Network Block-Cipher

Alex Biryukov, Technion
albi@cs.technion.ac.il

In Eurocrypt'91 Portz suggested the use of an interconnection network for the purpose of constructing a private key block cipher. Another suggestion for constructing ciphers through interconnection networks was made by Even and Yacobi. It turns out that neither of these systems is secure, although we believe this approach may lead to the construction ofsecure and fast ciphers yet.

A probabilistic polynomial-complexity cryptanalysis of the Portz cipher is presented. The analysis applies to a class of similar systems as well, and relies on the symmetry of the interconnection network and on the linearity of the control function for setting the switches. We outline possible ways to circumvent this weakness.

This is joint work with Shimon Even.

Early version: Technical Report CS0887, http://www.cs.technion.ac.il/Reports/

Cryptanalysis of RC5

Alex Biryukov, Technion
albi@cs.technion.ac.il

RC5 is a fast block cipher designed by Ron Rivest in 1994. Since then two attempts of cryptanalysis of this cipher were published. The best previously known attack requires 2^{54} chosen plaintexts in order to derive the full set of 25 subkeys for the 12 round RC5 with 32 bit words. In this paper we show a drastic improvement of these results. Our attack requires

2^{44} chosen plaintexts. We show that the 64 bit word version of RC5 is also much weaker than it was expected.

This is joint work with Eyal Kushilevitz, Submitted for publication. Early version: Technical Report CS0918, `http://www.cs.technion.ac.il/Reports/`

Comparing RSA and RSA-type Cryptosystems Over Elliptic Curves

Daniel Bleichenbacher, Bell Labs
bleichen@research.bell-labs.com
http://www.bell-labs.com/user/bleichen/

This talk compares RSA and RSA-type cryptosystems over elliptic curves. The authors of these elliptic curve cryptosystems hoped that their cryptosystems are more secure in broadcast applications (Hastad attack) in similar situations. Some new attacks are presented, which suggest that this hope is not justified. In particular, it is shown that some weaknesses of RSA when used with a small public exponent e can be extended to the discussed cryptosystems, but with no restriction on the public exponent e. This talk is mostly included in [1]. This paper can also be found on `http://www.bell-labs.com/user/bleichen/KMOV.ps`

[1] D. Bleichenbacher, "On the Security of the KMOV Public Key Cryptosystem", Advances in Cryptology - CRYPTO '97, vol. 1294. Springer-Verlag, pages: 235-248, 1997.

Twenty Years of Attacking RSA

Dan Boneh, Bellcore
dabo@cs.stanford.edu, http://www.cs.princeton.edu/~dabo/

The talk will give a survey of several attacks against the RSA system. These attacks can be classified into two categories: attacking the underlying mathematical structure, and the other attacking the implementation of RSA. The talk will focus on the beautiful mathematics involved in the first type of attacks. These attacks demonstrate the many pitfalls that security engineers must take into account in their designs.

Elliptic Curves

Len Charlap, IDA
len@ccr-p.ida.org

Before 1987, Elliptic Curves were thought to be just another interesting abstraction of pure mathematics. Then Hendrik Lenstra showed how you could use them to factor integers, a problem of practical value. Later, Victor Miller and Neal Koblitz pointed out that they could be used to strengthen important cryptological systems.

My talk will start with a brief explanation of Miller's idea and then give the basic theory of elliptic curves over finite fields. The main theme will be that most results in this area have elementary proofs that are usually more illuminating and useful (from a computational standpoint) than the proofs which use advanced ideas from Algebraic Geometry and Algebraic Number Theory. I will end with a heuristic explanation of why elliptic curves turn out to be so useful.

Computational Complexity Theory as an Applied Science

Joan Feigenbaum, AT&T Labs - Research
`jf@research.att.com, http://www.research.att.com/~jf`

Computational Complexity Theory is the study of "efficient" computation. Computational resources that one would like to use efficiently include time, space, randomness, and communication bandwidth. To prove that a solution to a computational problem is efficient, or alternatively that the problem has no solution that is efficient, one needs a formal model of computation, a well-defined notion of efficiency, and (sometimes) a formal way of proving that one problem is at least as "hard" as another. As a field, computational complexity has provided some of the most interesting and challenging open questions now faced by the mathematical world, the most famous of which is the notorious P vs. NP question.

An equally compelling proof of the importance of the field is its contribution to cryptology and security. The language and tools of computational compelxity make it possible for cryptologists and security experts to talk about crucial real-world concepts such as "intractability", "proof", "knowledge", "commitment", and "sharing" in a rigorous fashion. In this talk, I will review some of the fundamentals of computational complexity theory and argue that the theory is a critcial "enabler" for successful cryptology and security research.

The Radon Transform on the Hypercube

Ron Graham, AT&T Labs - Research
`rlg@research.att.com`

There are many situations in which information concerning a real-valued function f defined on a set X is available only in the form of averages of its values over various subsets of X. Such averages can be viewed as defining a transform F, called the Radon transform, of the function f. A fundamental question which arises is whether one can recover (in principle) f from knowledge of F.

In this talk (which is joint work with Persi Diaconis) we discuss this problem for several special classes of X, and, in particular, the n-dimensional hypercube, and show how these problems lead naturally to questions involving coding theory, NP-completeness, integer points on elliptic curves and linear recurrences of polynomials.

Recent Developments in Primality Testing

Jon Grantham, University of Georgia
`grantham@math.uga.edu`

The first half of the talk will be a brief exposition of developments in the past 15 years in the area of primality testing, including probable prime testing, primality proving using elliptic curves, and other techniqes of primality testing.

In the second half of the talk I'll describe a new probable prime test that takes three times as long as the so-called Miller-Rabin test, but is, in a certain sense, more than three times as accurate. I will also describe a result showing that there are general classes of probable prime tests which can be proved to have infinitely many pseudoprimes (composites which pass).

Papers are available on the World Wide Web at
`http://www.clark.net/pub/grantham/pseudo`

Cognitive Tutors: An Effective Technology to Improve Mathematics Learning

Ken Koedinger, Carnegie Mellon University
`koedinger@cmu.edu`

Student learning with the aid of an individual human tutor is dramatically better (2 standard deviations) than learning in traditional classrooms. A psychological theory of cognition and learning, John Anderson's ACT theory, provides an explanation for this effect and the basis for artificially intelligent learning environments called "Cognitive Tutors" that provide many of the positive features of human tutoring. Experiments with Cognitive Tutors have demonstrated dramatic gains in student learning in comparison with traditional instruction in mathematics and computer programming. Currently, a Cognitive Tutor for high school algebra is in use by students two days a week at over 25 schools. We are currently engaged in a cognitive analysis of the thinking processes involved in the subsequent high school mathematics courses, geometry and algebra II, and in the development and testing of Cognitive Tutors for these courses. We hope to demonstrate that three years of Cognitive Tutor use can lead to dramatic increases in students' mathematics achievement both in terms of practical problem solving and in academic preparation. Our prior results suggest one practical consequence of this effort is the raising of average math SAT scores by as much as 100 points. For further information, consult the web site `http://act.psy.cmu.edu/ACT/papers/koedinger-papers.html` or request copies of related articles from me at the email address `koedinger@cmu.edu`

Finding Short Vectors in Lattices

Jeff Lagarias, AT&T Labs - Research
`jcl@research.att.com`

This talk describes the problem of finding short vectors in lattices. It discusses what is known about the complexity of finding short vectors and presents the L^3 lattice basis reduction algorithm. It describes applications of this algorithm to breaking knapsack-type public key cryptosystems.

Cryptology, Technology and Policy

Susan Landau, University of Massachusetts at Amherst
`landau@cs.umass.edu`, `http://www.cs.umass.edu/~landau`

On April 16, 1993, the White House announced the Escrowed Encryption Standard (EES) – the Clipper chip, and a key escrow scheme with encryption keys split and stored with the government. The response to Clipper was substantial and negative. Despite that, the National Institute of Standards and Technology approved EES as a voluntary Federal standard for encryption of voice, fax, and computer information transmitted over circuit-switched telephone systems. The Clipper announcement presaged a four-year fight between the government and the public over the deployment of strong cryptography in the public sector, a fight that is continuing.

The shift to electronic communication – fax, email, data – has left society vulnerable to electronic eavesdropping. At the same time, law enforcement and national security will lose access to criminal communications if strong cryptography is easily available. What would be the cost to society if criminals concealed their communications using codes the government cannot decipher? How will U.S. economic competitiveness be affected by export controls on cryptographic systems? How important is protecting society from abuses by criminals and terrorists versus protecting personal privacy from all threats – including potential eavesdropping by the government?

In this talk I consider the dual-edged sword cryptography presents to both law enforcement and national security. I will present the debate on the deployment of cryptography in the context of related history and current needs.

This talk is based on the forthcoming "privacy on the Line: the Politics of Wiretapping and Encryption" by Whitfield Diffie and Susan Landau, MIT Press, February 1998.

`http://mitpress.mit.edu/book-home.tcl?isbn=0262041677`

When Can We Compute Square Roots?

Scott C. Lindhurst, University of Wisconsin
`lindhurs@math.wisc.edu`

In 1969, Daniel Shanks published an efficient algorithm for computing square roots mod p, that is, solving $x^2 = a(\text{mod } p)$. Shanks noted that his algorithm will, in fact, compute square roots in any cyclic group.

We show that a slight generalization of Shank's algorithm can compute square roots in many other groups, and characterize the groups in which the algorithm works.

Details can be found in the author's doctoral dissertation, "Computing Roots in Finite Fields and Groups, with a Jaunt through Sums of Digits" (University of Wisconsin-Madison, 1997).

Cryptographic Number Theory - Ignorance Is Bliss

Kevin McCurley, IBM - Research
`mccurley@almaden.ibm.com`

The White House recently released a document titled "A Framework for Global Electronic Commerce" in which it was predicted that electronic commerce on the Internet will amount to hundreds of billions of dollars within a very few years. Much of the security of the infrastructure for this commerce is ultimately based on some very fragile mathematical assumptions about our ignorance (e.g., that we can't efficiently find factors of large integers.) This talk will survey some of the most popular cryptographic assumptions in number-theoretic complexity and how they relate to each other. In particular, I will discuss some approaches to reducing the national security threat from innovative mathematicians. You'll hear many more unsolved problems than solutions.

Data Compression, Information, and Probability

Victor Miller, IDA
victor@ccr-p.ida.org

Programs that compress data have become ubiquitous these days. Even the very modem that you use contains one. Some of the ideas of data compression will be discussed. This will lead naturally to the idea of the "information" contained in a string of symbols. Along the way, the notion of "probability" and its meaning arise naturally.

Discrete Logarithms and Their Cryptographic Significance

Andrew M. Odlyzko, AT&T Labs - Research
amo@research.att.com

Given an element g of a group, the discrete logarithm of an element u in the subgroup generated by g is an integer k such that u = gk. The well-known problem of computing discrete logarithms in finite fields has attracted heightened increased attention in the last 20 years because of its applicability in cryptography. Related problems, such as that of the discrete logarithm problem on elliptic curves, have also become prominent recently. Security of many public key cryptosystems depend on the presumed difficulty of solving the discrete log problem. This lecture will survey this area.
http://www.research.att.com/~amo

Alternative Approaches to Integer Factoring

Rene Peralta, University of Wisconsin at Milwaukee
peralta@lucifer.cs.uwm.edu, http://cs.uwm.edu/faculty/peralta

The talk will describe my current work on integer factorization. This includes factorizations in progress, alternative approaches being investigated, and reductions to an NP-hard combinatorial problem. The latter problem is not known to be uniformly hard, hence it seems worthwhile to test the performance of standard algorithmic techniques on actual instances arising from integer factorization.

Constructing Hash Functions Based on Block Ciphers

Bart Preneel, Kuleuven
bart.preneel@esat.kuleuven.ac.be, http://www.esat.kuleuven.ac.be/~preneel

Collision resistant hash functions play an important role in cryptography. In this talk we review the construction of cryptographic hash functions based on block ciphers. The first such constructions date back to the late seventies; for these hash functions the size of the hash result is the same as the block length of the block cipher. We then discuss various proposals to obtain hash results larger than the block size, and describe some generic attacks. We conclude by presenting a new approach based on quaternary error-correcting codes (joint work with Lars Knudsen).

The resulting hash functions are more efficient than existing constructions; under reasonable assumptions about the underlying block cipher, one obtains collision resistant compression functions.

Cryptanalysis as Puzzles

Jim Reeds, AT&T Labs - Research
reeds@research.att.com

Cryptanalysis, or code-breaking, is the unauthorized decoding of encrypted messages. Particular cryptanalytic problems usually involve a lot of problem-specific technicalities, but at the heart of all cryptanalysis is puzzle solving. Whether the technical details of the particular problem obscure its puzzle-like nature depends on the complexity of the encryption system that the cryptanalyst must defeat. In this talk, I will solve some simple cryptanalytic puzzles in a way that exposes the analytic principles involved in the attacks.

Fast Arithmetic in GF[2^{156}]

Rich Schroeppel, University of Arizona
rcs@cs.arizona.edu
http://www.cs.arizona.edu/xkernel/www/people/rich.html

Using a tower-of-fields representation, arithmetic operations in GF[2^{156}] are three times as fast as our previous work with GF[2^{155}]. The case for elliptic curve cryptography is even more compelling.

Experimental Results on Efficient Generation of Shared RSA Keys

Sara Spalding, Indiana University and Rebecca Wright, AT&T Labs - Research
sspaldin@indiana.edu, rwright@research.att.com
http://www.research.att.com/info/rwright

We will first present the two-party version of Boneh and Franklin's protocol for efficient generation of shared RSA keys. The protocol allows two parties Alice and Bob to jointly compute an RSA key such that at the end of the protocol, the public key is known to both parties, and each party holds a share of the private key. Using their shares, Alice and Bob can cooperate to decrypt messages encrypted with the public key, but neither Alice nor Bob can decrypt messages alone. We will then present experimental results on the performance and optimizations of the protocol.

Review of Error Detection and Correction

Martin Strauss, AT&T Labs - Research
`mstrauss@research.att.com, http://www.research.att.com/~mstrauss`

We review the basics of error-detecting and -correcting codes, including models and definitions. Topics include linear codes (using Hamming codes as an extended example), how to form new codes from old (e.g., parity check and concatenation), and bounds on the sizes of codes (e.g., sphere-packing, linear programming, and Gilbert-Varshamov). We also review some connections between error-correction and cryptology, including secret-sharing, hardness of decoding, and the relationship between encryption and error-correction in real-world channels.

Space Efficient Group Structure Computation Using Pollard's Rho-Method

Edlyn Teske, University of Manitoba
`teske@cs.umanitoba.ca`
`http://www.informatik.th-darmstadt.de/TI/Mitarbeiter/teske.html`

We present a new algorithm for computing the structure of a finite abelian group. This algorithm is based on Pollard's Rho-Method for index computation and has to store only a fixed, small number of group elements, independent of the group order. It is generic in the sense that it does not exploit very special properties of the group operations or the encodings of the group elements.

We estimate the computational complexity of our algorithm by counting the group operations such as multiplications and equality checks. Under some plausible assumptions, we prove that the expected run time is $O(\sqrt{n})$ (with n denoting the group order), and we explicitly determine the O-constants.

We implemented our algorithm for ideal class groups of imaginary quadratic orders and present a selection of our experimental results.

Edlyn Teske, "A space efficient algorithm for group structure computation", to appear in Mathematics of Computation, 1998. See also Technical Report TI-3/97, Technische Universitat Darmstadt, Germany.

Available via http://www.informatik.tu-darmstadt.de/TI/Mitarbeiter/teske.html

On Provable Security for Secret Key Block Ciphers

Serge Vaudenay, ENS
Serge.Vaudenay@ens.fr
http://www.dmi.ens.fr/dmi/users/vaudenay/index_en.html

Since the Data Encryption Standard has been proposed, the security of block ciphers relies on heuristic and empirical arguments. This situation has a changed a little bit since Biham and Shamir's differential cryptanalysis breakdown and a class of statistical attacks emerged against block ciphers, as well as design criteria for thwarting them.

In this talk we investigate methods for providing provable security against a wide class of attacks. We show that a very simple (and hopefully cheap) combinatorial concept can make them secure and raise a bridge between Shannon's perfect secrecy concept and the empirical approach for security.

Speeding Up Public Key Systems via Pre-Computations

Ramarathnam Venkatesan, Microsoft Research
venkie@microsoft.com

We present practical schemes for speeding up public key schemes based on factoring and discrete log. To achieve this, we generate certain distributions on the instances of the form (x, g^x) or (x, x^e) and analyze the security of the resultant systems. Analyzing what realistic attacks would take leads to some hidden lattice problems, which may be of independent interest.

Joint work with V. Boyko (MIT) and M. Peindao (GMD Research).

Comparing Without Revealing

Peter Winkler, Bell Labs
pw@research.bell-labs.com, http://www.research.bell-labs.com/~pw

The following problem arises in a variety of circumstances (e.g. bargaining, gossiping, passwords, entrapment): two people each possess some information and wish to discover if the information is the same, but without revealing any of the information in the case that it is not. How can they do it? In particular, how can they do it if they are not both experts in public key cryptography?

We have collected a number of solutions, involving everything from number theory to playing cards, paper cups and airline reservations. Not all the approaches are perfect probabilistically or applicable in all cases, but all can actually be used in practice. We invite you to contribut further ideas to our list.

Joint work (more accurately, joint fun) with Ron Fagin (IBM) and Moni Naor (The Weizmann Institute).

Secret Communication Using a Deck of Cards

Rebecca Wright, AT&T Labs - Research
`rwright@research.att.com, http://www.research.att.com/info/rwright`

In card games such as Bridge, partners try to communicate information about their cards to each other while concealing information from the opposing partners. Based on this idea, we will investigate the use of an ordinary deck of cards to communicate secrets that are _completely_unrelated_ to the actual cards. We will start by looking at a setting with three players: Alice, Bob, and an eavesdropper, Eve. Alice wishes to tell Bob the answer to a yes/no question while concealing the answer from Eve. We will explore methods Alice and Bob might use and the conditions needed to achieve secret communication. No Bridge experience necessary!

Week 2 (August 3 - August 9, 1997): Computer Science

Commodity-based Cryptography

Don Beaver, IBM/Transarc
beaver+@transarc.com

This work introduces a new paradigm for the design of protocols for secure joint computation requiring minimal interaction. Instead of relying on trusted and specialized devices, unproven cryptographic assumptions, or highly interactive multiparty computations, this work proposes a commodity-based model in which servers provide security resources to clients but are not involved in the clients' computations themselves. Restricting the involvement of servers in turn improves scalability, simplicity, and security.

Unlike oracles, which typically provide computational resources such as the results of infeasible computations, these servers assist clients in establishing shared resources for secure computations such as oblivious transfer and circuit evaluation, broadcast, and multiparty computations. Unlike protocols for secure multi-party computation, the servers themselves are "non-interactive", and in fact, have no knowledge whatsoever of each other. They provide security resources to each client through a single RPC.

We give explicit constructions supporting oblivious transfer and (time permitting) network multicast.

Concrete Cryptography

Don Beaver, IBM/Transarc
beaver+@transarc.com

Cryptography is often a matter of rediscovery as much as discovery: the most elegant of solutions has often already appeared as a children's game or a common social principle. As a result, explaining cryptography in simple and concrete terms, using simple and concrete props, is not just a means to convey an understanding of otherwise lackluster mathematical concepts but is also a useful tool for developing and discovering new cryptographic tools. The greater the complexity of a solution, the more likely it is insecure: thus a fourth-grade explanation is both a pedagogic tool and a screen for security.

We'll discuss straightforward and concrete implementations of some of the most fundamental tools in cryptography, including shared secrets, deniability, and "oblivious transfer", and we'll see where children have already beat cryptographers to the punch. The mathematical counterparts to these concrete solutions involve elementary modular arithmetic and sometimes require a childlike ability to think "sideways".

Reliable Communication Over Partially Authenticated Networks

Amos Beimel, DIMACS
beimel@dimacs.rutgers.edu
http://dimacs.rutgers.edu/People/Beimel.html
http://dimacs.rutgers.edu/~beimel

Reliable communication between parties in a network is a basic requirement for executing any protocol. In this talk, we consider the effect on reliable communication when some pairs of parties have common authentication keys. The pairs sharing keys define a natural "communication graph", which may be quite different from the "communication graph" of the network. We characterize when reliable communication is possible in terms of these two graphs, focusing on the very strong setting of a Byzantine adversary with unlimited computational resources.

Publication: A. Beimel, and M. Franklin. "Reliable communication over partially authenticated networks", in WDAG'97, volume 1320 of Lecture Notes in Computer Science. pages 245-259. Springer, 1997. Also, DIMACS technical report 97-41, available at http://dimacs.rutgers.edu/TechnicalReports/1997.html

The Impact of Quantum Mechanics on Cryptology

Gilles Brassard, Universit de Montral
brassard@iro.umontreal.ca

Quantum mechanics has the potential to play a major role in the future of cryptology. On the one hand, it could bring to its knees most of the current trends in contemporary cryptography. On the other hand, it offers an alternative for the protection of privacy whose security cannot be matched by classical means. I shall review both sides of the coin. No prior knowledge of quantum mechanics will be assumed.

An Overview of Unconditionally Secure Key Agreement

Christian Cachin - MIT
cachin@acm.org

All cryptosystems in use today can theoretically be broken by an exhaustive search of the key space. In this talk, we will describe how information-theoretic methods can be used for proving the security of unconditionally secure cryptosystems, with the focus on key agreement protocols between participants that do initially not share secret information.

The operation of many such systems can be divided into three phases called advantage distillation, information reconciliation, and privacy amplification. We illustrate some of the protocols used and discuss related questions. As one specific example, we also present a key agreement protocol that is unconditionally secure based on the sole assumption that an adversary's memory capacity is limited.

Further information on this subject can be found in: Christian Cachin. Entropy Measures and Unconditional Security in Cryptography, volume 1 of ETH Series in Information Security and Cryptography. Hartung-Gorre Verlag, Konstanz, Germany, 1997. ISBN 3-89649-185-7 (Reprint of Ph.D. dissertation No. 12187, ETH Zurich).

Christian Cachin and Ueli Maurer. Unconditional security against memory-bounded adversaries. In Burt Kaliski, editor, Advances in Cryptology: CRYPTO '97, volume 1294 of Lecture Notes in Computer Science, pages 292-306. Springer-Verlag, 1997.

Towards Realizing Random Oracles: Hash Functions That Hide All Partial Information

Ran Canetti, IBM - Research
canetti@watson.ibm.com, http://theory.lcs.mit.edu:80/~canetti/

The random oracle model is a very convenient setting for designing cryptographic protocols. However, we do not know how to construct a mechanism that realizes a random oracle. In fact, we do not even know how to meaningfully *specify* the properties required from such a mechanism. We initiate an effort to improve this situation, by proposing a new primitive that realizes a specific aspect of random oracles. This primitive, called *oracle hashing*, is a hash function that, like random oracles, 'hides all partial information on its input'. A salient property of oracle hashing is that it is probabilistic: Different applications to the same input result in different hash values. Still, we maintain the ability to verify whether a given hash value was generated from a given input. We show several applications and constructions of the new primitive.

An Extended Abstract of this work appears in the proceedings of CRYPTO 97.

On the Adaptive Security of Multiparty Protocols

Ran Canetti, IBM - Research
canetti@watson.ibm.com, http://theory.lcs.mit.edu:80/~canetti/

A fundamental problem in the design of secure multiparty protocols is how to deal with adaptive adversaries (i.e., adversaries that may choose the corrupted parties during the course of the computation.) This problem brings forth concerns that were not addressed otherwise. For instance, the power of an adaptive adversary is greatly affected by the extent to which (even uncorrupted) parties are trusted to carry out instructions that cannot be externally verified, such as erasing history records of the execution. We describe the importance of adaptive security, its definition, and survey recent constructions of adaptively secure protocols. We also point out some directions for further research.

Much of the above talk is covered in a paper with Uri Feige, Oded Goldreich and Moni Naor, avaliable as MIT LCS TR 682, 1996 (extended abstract at STOC '96).

Private Information Retrieval

Benny Chor, Technion
`benny@cs.technion.ac.il`

Publicly accessible databases are an indispensable resource for retrieving up-to-date information. But accessing such databases also poses a significant risk to the privacy of the user, since a curious database operator can follow the user's queries and infer what the user is after. A trivial (but totally impractical) solution to the private information retrieval (PIR) problem is to download the whole database. The question is whether such privacy can be achieved at more reasonable communication costs.

In this survey talk we will describe several solutions to the PIR problem. Most of these solutions rely on replicating the database among k ($k \geq 2$) non-communicating servers. The more expensive solutions (in terms of communication complexity) guarantees information theoretic privacy. Substantially less expensive solutions exist if one is ready to settle for computational privacy.

The Future of Privacy

Lorrie Cranor, AT&T Labs - Research
`lorrie@research.att.com`

Online privacy concerns have been the focus of increasing amounts of attention from the media, legislators, and others in the United States and abroad. Individuals are concerned about the ability of Web sites to obtain information about them without their knowledge, the widespread availability of their personal information in networked databases, and the use of their personal information to send them unwanted solicitations (including junk email, commonly referred to as "spam"). A variety of solutions have been proposed to address each of these concerns, including new laws and regulations, industry self-regulation, and user empowerment technologies that automate individual control. In this talk I will describe several online privacy concerns and discuss possible solutions that may address these concerns.

For further information please see: The Role of Technology in Self-Regulatory Privacy Regimes. In Privacy and Self Regulation in the Information Age. U.S. Department of Commerce, National Telecommunications and Infrastructure Administration. June 1997. p. 185-191. `http://www.research.att.com/~lorrie/pubs/NTIA.html`

Cryptographic Power of Noisy Channel

Claude Crépeau, Université de Montréal
`crepeau@iro.umontreal.ca`, `http://www.iro.umontreal.ca/~crepeau/`

We consider several cryptographic scenarios of two or three people attempting to achieve information theoretic protocols. We consider such scenarios as the Key Distribution protocol or the Secure Function Evaluation protocol and show how these can be implemented securly based on the existence of a noisy channel generating errors on each transmitted bit with a fixed probability.

Discreet Solitary Games

Claude Crépeau, Université de Montréal
`crepeau@iro.umontreal.ca`, `http://www.iro.umontreal.ca/~crepeau/`

It's nearly Christmas time and you have to buy presents for your family and friends. Indeed, among certain families there is a more economical approach to this situation: rather than buying one present per person, each member of a group picks the name of another member and becomes responsible for buying that person a present.

Traditionally, the one person for whom each member is responsible, is allocated at random using the "names-in-a-hat" technique: each person puts his/her name in a common hat and then everybody picks a name at random from the hat. If by accident one picks one's own name, he/she puts it back. Otherwise everyone is responsible for the present of the person he/she picked.

To put it abstractly, the goal is for the n persons involved to pick a random permutation $\pi(i)$. Now consider the scenario where the members of this family cannot be gathered in a room to do the "names-in-a-hat" technique, for instance if some of them live abroad in several different countries.

Cryptographic techniques have been used intensively in the past to show how to play multiparty games in an adversarial scenario. We now investigate the cryptographic power of a deck of cards in a solitary scenario. In particular, we show how a person can select a random permutation satisfying a certain criterion (discreetly - without knowing which one was picked) using a simple deck of cards. We also show how it is possible using cards to play games of partial information such as POKER, BRIDGE and other card games in isolation.

Does It Matter What Students Actually Think?

Robert Davis, Rutgers University
`RDavisRU@aol.com`

In an age when it is more important than ever for mathematicians and mathematics educators to work closely together, we need to understand one another. As one report from math ed studies, we will look at evidence of how students actually think. One thing, at least, is clear - they do not necessarily think the thoughts we want them to!

Positive Applications of Lattices to Cryptography

Cynthia Dwork, IBM - Research
`dwork@almaden.ibm.com`

Initiated by Ajtai's paper "Generating Hard Instancesof Lattice Problems", a burgeoning effort to build cryptographic primitives based on the assumed hardness of worst-case or random instances of problems involving lattices has proved extremely fruitful. Prior to Ajtai's work, lattices, and in particular, the lattice basis reduction algorithm of Lenstra, Lenstra, and Lovasz, were used in cryptography principally to prove cryptographic insecurity. More positive applications of lattices include schemes for identification, bit commitment, coin flipping, public key cryptosystems, cryptographically strong hash functions, and pseudo-random number generators, each of whose security depends only on the worst-case hardness of the underlying lattice problem; as well as a digital signature scheme whose security depends on the average hardness of the underlying problem. The talk will discuss some of these constructions.

Random-self-reducibility and Instance-hiding: Overview, Applications, and Open Questions

Joan Feigenbaum, AT&T Labs - Research
`jf@research.att.com, http://www.research.att.com/~jf`

A function of f is "random-self-reducible" if the evaluation of f at any given instance x can be reduced in polynomial time to the evaluation of f at one or more random instances

y_i. Random-self-reducible functions have long found applications throughout the computer-science research world, e.g., in average-case complexity, lower bounds, interactive proof systems, secure protocol design, and program checking, testing, and self-correcting. Active areas of current research to which random-self-reducibility is closely related include private information retrieval, lattice-based cryptosystems, and cryptographic data structures.

Despite many years of notable progress, random-self-reducible functions are still far from completely classified. More generally, many interesting open questions remain about random-self-reducibility per se and about its relationship to other fundamental notions in cryptology and complexity. In this talk, I will review the basics of this topic, explain its relationship to some currently hot topics, and state some of the basic open questions.

Kid Crypto

Mike Fellows, University of Victoria
`mfellows@imada.ou.dk`

Cryptography is now a very broad area of mathematical research that has many exciting, dramatic and important applications. The effort to make this area of contemporary mathematical science accessible to students in grades K-12 (at varying depths of sophistication, of course) has come to be called "kid krypto". Serious "industrial strength" cryptography is frequently modular, in the sense that there are basic building blocks (such as "bit commitment") that are used to assemble more complicated information-exchange protocols. At the most basic level, mathematical tools and concepts from areas such as number theory, probability theory and complexity theory usually provide the most basic building blocks. The modularity of cryptography lends itself to a kind of playful deconstruction where some of the more sophisticated and demanding mathematical components can be replaced by simpler basic parts that are more accessible to a K-12 audience (with the resulting systems being less secure or efficient, but still interesting).

The result is that some of the remarkable "gems" of cryptography, such as coin flip protocols and public key systems, can be engaged using only such ingredients as elementary combinatorics, arithmetic, boolean logic, and polynomials. Cryptography can thus serve as a source of some contemporary excitement and as a window on the work of research mathematicians and computer scientists by building on these standard curriculum topics. Kid krypto is by no means a finished subject – there are a variety of possibilities for original research that could be undertaken by high school students or faculty.

New Techniques for Sharing Cryptographic Functions

Yair Frankel, CERTCO
`frankely@certco.com`, `http://www.cs.sandia.gov/~yair/`

We introduce new techniques for sharing cryptographic functions in a distributed manner using share-of-share protocols which we call: "poly-to-sum" and "sum-to-poly". Employing these techniques, we solve some open problems and develop new efficient protocols in the area of theshold cryptography (cryptographic function sharing). A new optimal resilience proactive threshold scheme and a new robust "share-size efficient" robust RSA function sharing protocol is developed.

Ecash and the Power of Positive Paranoia

Matthew Franklin, AT&T Labs - Research
franklin@research.att.com

The suspicions of ecash researchers extend to both bank robbers and banks, extortionists and charities, counterfeiters and governments. I will discuss a few of the elegant solutions that justify this creative mistrust.

Copyright 101 for Digital Domain

Brian A. LaMacchia, AT&T Labs - Research
bal@microsoft.com, http://www.farcaster.com/

The confluence of copyright law and digitally-encoded information presents, depending on whom you ask, either a potential boon or disaster (or both!) for those involved in the "intellectual property" business. This talk will provide a brief history of the development of U.S. copyright law and then use that context to discuss some of the current issues involving the protection and distribution of intellectual property over digital networks. In addition, we will present some proposed technological systems for enumerating, controlling, and enforcing in the digital domain the "bundles of rights" held by content creators, publishers, consumers and the public. We will also look at how such systems, including digital rights management systems, would interact with proposed legislative changes in copyright law.

On-line copies of the slides presented during this talk are available from the author's homepage at `http://www.farcaster.com/`

Number-Theoretic Constructions of Efficient Pseudo-random Functions and Other Cryptographic Primitives

Moni Naor, The Weizmann Institute
naor@wisdom.weizmann.ac.il

We describe efficient constructions for various cryptographic primitives (both in private-key and in public-key cryptography), based on the decisional version of the Diffie-Hellman assumption and on the hardness of factoring.

Our major result is a new and efficient construction of pseudo-random functions. Computing the value of the function at any given point involves two multiple products (which is comparable with two exponentiations). In particular, the functions are shallow, they can be computed in TC^0 (the class of functions computable by constant depth circuits consisting of a polynomial number of threshold gates).

Using the simple algebraic structure of the pseudo-random function, f_s, we show a zero-knowledge proof for statements of the form "$y = f_s(x)$" and additional features of the functions.

Joint work with Omer Reingold. A full paper is available at
`http://www.wisdom.weizmann.ac.il/~reingold/research.html`

Private and Secure Database Storage and Retrieval

Rafail Ostrovsky, Bellcore
`rafail@bellcore.com`

In this talk, communication-efficient ways to store and retrieve data from the database will be described. The talk will cover material from two papers: the paper from previous STOC, titled "Private Information Storage" (joint work with Victor Shoup) and the paper which is to appear in the forthcoming FOCS, titled "Replication Is Not Needed: A Single Database Computational Private Information Retrieval" (joint work with Eyal Kushilevitz). The talk will be self-contained.

Randomness-Efficient Non-Interactive Zero-Knowledge Proofs

Pino Persiano, Universita' di Salerno
`giuper@dia.unisa.it`

The model of Non-Interactive Zero-Knowledge allows one to obtain minimal interaction between prover and verifier in a zero-knowledge proof if a public random string is available to both parties. In this talk we investigate upper bounds for the length of the random string for proving one and many statements, showing the following results: We show how to prove in non-interactive perfect zero-knowledge any polynomial number of statements using a random string of fixed length, that is, not depending on the number of statements for the case of Quadratic Residuosity. Under the quadratic residuosity assumption, we show how to prove any NP statement in non-interactive zero-knowledge on a random string of length

$\Theta(nk)$, where n is the size of the statement and k is the security parameter, which improves the previous best construction by a factor of $\Theta(k)$.

Identity Escrow

Erez Petrank, DIMACS
`erez@dimacs.rutgers.edu`,
`http://dimacs.rutgers.edu/People/Postdocts/Petrank.html`

We introduce the notion of *escrowed identity*, an application of key-escrow ideas to the problem of identification. In escrowed identity, one party A does *not* give his identity to another party B, but rather gives him information that would allow an authorized third party E to determine A's identity. However, B receives a guarantee that E can indeed determine A's identity. We give protocols for escrowed identity based on the El-Gamal (signature and encryption) schemes and on the RSA function. A useful feature of our ptotocol is that after setting up A to use the system, E is only involved when it is actually needed to determine A's identity.

This is joint work with Joe Killian.

International patent pending, April '97.

DIMACS Technical Report 97-28,
`http://dimacs.rutgers.edu/TechnicalReports/1997.html`

Web access at The Theory of Cryptography Library, Item 97-11
`http://theory.lcs.mit.edu/~tcryptol/1997/97-11.html`

Chameleon Hashing and Signatures

Tal Rabin, IBM - Research
`talr@watson.ibm.com`

We introduce chameleon signatures that provide with an undeniable commitment of the signer to the contents of the signed document (as regular digital signatures do) but, at the same time, do not allow the recipient of the signature to disclose the contents of the signed information to any third party without the signer's consent.

These signatures are closely related to Chaum and van Antwerpen's undeniable signatures, yet chameleon signatures allow for simpler and more efficient realizations. In particular, they are essentially non-interactive and do not involve the design and complexity of zero-knowledge proofs on which traditional undeniable signatures are based. Instead, chameleon signatures are generated under the standard method of hash-then-sign. Yet, the hash functions which are used are chameleon hash functions. These hash functions are characterized by the non-standard property of being collision-resistant for the signer but collision tractable for the recipient.

We present simple and efficient constructions of chameleon hashing and chameleon signatures. The former can be constructed based on standard cryptographic assumptions (such as the hardness of factoring or discrete logarithms) and have efficient realizations based on these assumptions. For the signature part we can use any digital signature (such as RSA or DSS) and prove the unforgeability of the underlying digital signature in use.

Our schemes are non-interactive and do not involve the design and complexity of zero-knowledge proofs, which form the basis of traditional undeniable signatures.

The paper appears in `http://www.research.ibm.com/security/chameleon.ps`

Pseudorandomness Against Non-deterministic Adversaries

Steven Rudich, Carnegie Mellon University
`rudich@cs.cmu.edu, http://www.cs.cmu.edu/~rudich`

The notion of pseudorandom can be generalized to work against an adversary powerful enough to guess the seed. This seeming contradiction leads to several interesting open questions.

Cryptography Without Computers

Adi Shamir, The Weizmann Institute
`shamir@wisdom.weizmann.ac.il`

In the last 50 years, cryptology had become increasingly computerized. In this talk I'll describe several novel cryptographic and cryptanalytic techniques which can be carried out without using any computers. In particular, I'll introduce a new paradigm for carrying out massively parallel key searches, which is much simpler to implement than alernative approaches based on DNA molecules or quantum effects.

A New Paradigm for Massively Parallel Random Search

Adi Shamir, The Weizmann Institute
`shamir@wisdom.weizmann.ac.il`

The problem of optimizing combinatorial problems or breaking cryptographic codes led to several novel paradigms for carrying out such a massively parallel random search, including quantum and DNA computers. In this talk I'll propose a new paradigm, which is based on a simple and easy to implement idea.

Proxy Cryptography

Martin Strauss, AT&T Labs - Research
`mstrauss@research.att.com`, `http://www.research.att.com/~mstrauss`

We introduce *proxy cryptography*, in which a *proxy function*, in conjunction with a public *proxy key*, converts ciphertext (messages in a public key encryption scheme or signatures in a digital signature scheme) for one key k_1 into ciphertext for another k_2. Proxy keys, once generated, may be made public and proxy functions might exist; *symmetric* proxy functions assume that the holder of k_2 unconditionally trusts the holder of k_1, while *asymmetric* proxy functions do not. We give examples of proxy schemes for encryption and signature schemes.

This is joint work with Matt Blaze.

Combinatorial Cryptology and the 'Two Sheriffs' Problem

Peter Winkler, Bell Labs
`pw@research.bell-labs.com`, `http://www.research.bell-labs.com/~pw`

Two sheriffs in neighboring counties are looking for the killer in a case involving eight suspects. By reliable detective work, each sheriff has independently narrowed the list to only two; now they're on the phone together, and wish to compare lists so that (if they haven't narrowed down to the same two suspects) they can deduce the identity of the killer and capture him, no matter which country he's in.

The difficulty is that the local lynch mob has tapped the phone, and if they can deduce the identity of the killer he will be lynched before the sheriffs can bring him in.

The sheriffs are speaking to one another for the first time, and have no protocol prepared in advance. Can they nonetheless conduct a conversation which will result in both knowing the killer (when possible) while leaving the lynch mob in doubt?

We will provide the answer, the theory, and two more equally dubious applications. Joint work with Don Beaver (IBM/Transarc) and Stuart Haber (Surety Technologies).

Week 3 (August 10 - August 15, 1997): Network Security

Strengthening Passwords

Martin Abadi, Digital Systems Research Center, DEC
ma@pa.dec.com
http://www.research.digital.com/SRC/personal/Martin_Abadi/home.html

Despite much progress in cryptographic authentication, and despite their notorious vulnerability, traditional passwords are likely to remain important for security. This talk discusses a method for strengthening traditional passwords. The method does not require users to memorize or to write down long passwords, and does not rely on smart-cards or other auxillary hardware. The main cost of the method is that it lengthens the process of checking a password.

This is joint work with Mark Lomas and Roger Needham.

Security in Clinical Information Systems

Ross Anderson, Cambridge University
ross.anderson@cl.cam.ac.uk, http://www.cl.cam.ac.uk/users/rja14/

Over the last two years, there has been sharp controversy in the UK between the government, which has been attempting to build a national healthcare network, and the medical profession, which opposes the centralisation of personal health information on both safety and privacy grounds. The fundamental tension is that government administrators believe that they have the right to take access control decisions based on 'need to know', whereas medical ethics states unambiguously that control rests with the patient: the patient has an almost unfettered right to forbid third parties such as bureaucrats from accessing his medical record.

This controversy has opened up a number of fascinating new research topics. How does one build a security system in which access rights are not determined by a central administrator, but by the data subjects themselves? How can trust be managed economically in such a system? What sort of security policy models are appropriate? Is key escrow workable at all in large scale distributed systems?

These questions have importance far beyond medicine. Essentially the same trust relationships apply in other professions, and so the lessons learned from medicine may be directly useful to lawyers and accountants (among many others). In fact, our medical work is a good first step towards a general protection policy model for personal privacy in a very broad sense.

Steganography

Ross Anderson, Cambridge University
`ross.anderson@cl.cam.ac.uk`, `http://www.cl.cam.ac.uk/users/rja14/`

Cryptography enables us to conceal the content of a message, but sometimes we need to do more - we may have to hide the identity or physical location of the sender, the receiver or both. In some applications, we need to conceal the message's very existence.

Steganography is the art of hiding information in other information. It has recently attracted much attention from the research community, driven by the desire of intellectual property owners to hide encrypted copyright marks and serial numbers in digital audio and video works, and the proposals by a number of governments to restrict the use of cryptography.

In this talk, I will provide a historical overview of the subject, look at the main contending technologies, describe some novel attacks, and discuss the prospects of developing a theory of the subject.

Information Theoretic Key Distribution Schemes

Amos Beimel, DIMACS
`beimel@dimacs.rutgers.edu`
`http://dimacs.rutgers.edu/People/Beimel.html`
`http://dimacs.rutgers.edu/~beimel`

Key management plays a fundamental role in cryptography as a basis for securing cryptographic techniques providing confidentiality, entity authentication, data integrity and digital signatures. In this talk we will discuss Key Distribution Schemes in which during an initialization stage a trusted server generates and distributes secret data pieces to the users, such that subsets of the users may subsequently compute a secret shared key. The security of the schemes we will consider is not based on any assumptions (i.e., information theoretic secure).

Analysis of the SSL Protocol

Sven Dietrich, Adelphi University
`spock@abraxas.adelphi.edu`

We present an overview of formal methods for the specification and analysis of authentication protocols and a formal specification and analysis of the Secure Sockets Layers Protocol v3.0. The protocol is specified using an extension of the Rubin-Honeyman NCP logic, based on knowledge and belief, developed for nonmonotonic cryptographic protocols. The existing logic is extended to fit the purpose and the analysis is performed for three specific cases. We draw conclusions about the assumptions of the protocol in general, show the weak points of the protocol, and outline possible attack techniques.

Further research work may be found at `http://www.adelphi.edu/~spock`

Working Cryptanalysis of the German Enigma

Carl Ellison, Cybercash
`cme@cybercash.com`

The 3-rotor Enigma used by the majority of the German armed forces in WW-II had a keyspace of 67.6 bits. Current practice in 1997 allows us to build a custom computer to test in excess of 200 million keys per second for each $10 component. A machine costing one million dollars could then test twenty trillion keys per second. That machine could search the Enigma keyspace in 126 days. Assuming that computing power grows by a factor of 2 every 1.5 years, that machine in 1940 technology would have required 95 billion years to do the same test.

Thanks to the algebraic structure of the cipher and the genius of Alan Turing and Gordon Welchman, the British Bombe designed in 1940 did a key search in a maximum of 15 hours. Because the problem could be parallelized, the minimum key search time was 6 minutes (after manual setup). This talk describes the operation of the Enigma machine and the Turing/Welchman Bombe, including an audience simulation of the Bombe in action.

Key Management in the Post-Identity Era

Carl Ellison, Cybercash
`cme@cybercash.com`

Key management is a topic as old as cryptography itself. With the invention of public-key cryptography, it was presumably made simpler. With the invention of public-key identity certificates, binding names of people to their keys, it was made simpler still (in theory). The Global Village, especially in the form of the Internet, has created a perfect arena for the employment of modern key management – but it has brought with it two unanticipated consequences which, in turn, have made key management nearly as difficult as it ever was. The first of those two is that, with so many people, the traditional notion of “identity” (as a person's name) has been destroyed. This talk will cover the history of key management, through the age of identity certificates, then describe the death of identity certifications and discuss the emerging needs of a key-management system in the post-identity era.

`http://www.clark.net/pub/cme/html/spki.html`

How to Sign Digital Streams

Rosario Gennaro, IBM - Research
`rosario@watson.ibm.com`, `http://theory.lcs.mit.edu:80/~rosario/`

We present a new efficient signature paradigm to sign digital streams. The problem of signing digital streams in order to prove their authenticity is substantially different from the problem of signing regular messages. Indeed, traditional signatures assume that the receiver holds the entire message being signed in order to authenticate it. However, a stream is a potentially very long (infinite) sequence of bits that the sender sends to the receiver. Moreover, the receiver consumes the data it receives at more or less the input rate. This means that it is infeasible for the receiver to obtain the entire stream before authenticating it. Examples of streams include digitized audio and video files, and applets.

We present two solutions to the problem of authenticating digital streams. The first one is for the case of a finite stream which is entirely known to the sender (say a movie). We use this constraint to devise an extremely efficient solution. The second case is for a (potentially infinite) stream which is not known in advance to the sender (for example a live broadcast). We present proofs of security of our constructions.

Our techniques have also applications in other areas as for example efficient authentication of long files when communication is at cost.

This is joint work with P. Rohatgi.

The full paper is available from `http://www.research.ibm.com/security/publ.html`

Ensuring the Integrity of Records On Line: How to Time-Stamp a Digital Document

Stuart Haber, Surety Technologies
`stuart@surety.com`

Encryption can be used to ensure the privacy of electronic records, and digital signatures can be used to identify the author of a record. But, until recently, there was no satisfactory answer to the following question: How can one know with certainty when a digital document was created or last modified, and that it has not been altered by anyone, including its author, since that time? This problem has gained in importance as more and more of the world's records are created, manipulated, transmitted, and stored entirely in digital form.

This talk presents cryptographically secure digital time-stamping, a solution to this problem developed by the speaker and Scott Stornetta. Users of their Digital Notary[tm] system can *certify* their digital documents, computing for any particular document a concise time-stamp *certificate*. Later, any user of the system can *validate* a document-certificate pair, verifying that the document existed in exactly its current form at the time asserted in the certificate. The security of the system depends on the use of a cryptographic tool called a one-way hash function; any adversary wishing to compute a false or back-dated certificate must break the hash function in order to succeed in doing so. Surprisingly, there is no requirement that an agent be trusted or that a cryptographic key be kept secret.

Another problem raised by the widespread use of computers is that of "naming" digital documents in such a way that any user in possession of a document can be sure that it is indeed the one that is referred to by its name. The problem is especially acute on the World-Wide Web, where a document (whose only existence may be on line) is now typically named by giving its URL, which is merely a pointer to its virtual location at a particular moment in time.

With a simple variation of the time-stamping process, the Digital Notary system can also provide a cryptographically secure name or serial number for any certified document, one that depends on all the bits in the document but is only about the length of a telephone number. The secure link between a document and its "name" can be validated by a variant of the usual validation procedure for time-stamp certificates.

Finally, time-stamping can be used in certain circumstances to extend the useful lifetime of different sorts of cryptographic certifications of authenticity, in the event that the certifying protocol is compromised. This can be applied to digital signatures, or to time-stamping itself, making the digital time-stamping process renewable.

A commercial implementation of the Digital Notary system is available from Surety Technologies, a spin-off of Bell Communications Research (see `http://www.surety.com`). The system can be used for any kind of computer file or electronic record whatsoever, including text, audio, video, drawings, images, formatted publications, spread-sheets, data-base entries, and e-mail messages.

Probability Theory for Pickpockets

Marcus Kuhn, Purdue University
`kuhn@cs.purdue.edu, http://wwwcip.informatik.uni-erlangen.de/~mskuhn`

Why would a card thief be interested in probability theory? We discuss how knowledge of conditional probabilities and Bayes' theorem can help a criminal to steal your money and how ignorance in probability theory can get you in trouble as a bank computer programmer for the very same reason.

We look at the PIN generation and verification procedure of the EuroCheque debit card, the magnetic stripe card that most Germans use to get their daily cash from automated teller machines. It turns out to be a nice example of how the wrong application of a high-quality encryption algorithm (DES) can result in a surprisingly insecure overall system.

All calculations use only high-school level probability theory and involve a simple computer program. This makes the PIN guessing problem a fascinating student project that can illustrate everyday computer-security mechanisms in a math class.

The paper is available on `http://www.cl.cam.ac.uk/~mgk25/pinprob.pdf`

Music on the Internet and the Intellectual Property Protection Problem

Jack Lacy, AT&T Labs - Research
`lacy@research.att.com`

Recent advances in audio compression technology, coupled with trends toward higher network bandwidth, lower memory costs, and lower storage costs, indicate that the obstacles to music distribution on the Internet will soon be overcome. Because it provides benefits to consumers and to music content owners alike, network music distribution may become one of the first compelling Internet consumer services. At the same time, these technological advances create opportunities for pirates. We believe that the music industry must prepare for network music distribution or face widespread theft of its music.

In this talk, we first review the technological advances that are driving a change in the industry's business mode. We then discuss the opportunities that these advances allow. Finally we focus on the dangers to content owners and discuss various ways to protect musical content.

This is joint work with David P. Maher and James H. Snyder.

Security Models for Partially Accounted E-Cash Systems

David P. Maher, AT&T Labs - Research
`dpm@research.att.com`

We discuss security models for e-cash systems where cash balances are distributed among many certified, secure microcomputers (purses) that can interact over networks. Changes in cash balances occur when two purses, after verifying each other's credentials, follow certain rules and exchange appropriate protocol messages

Two types of system break-ins can occur. The first involves the counterfeiting of new purses that appear (to other purses) to be properly certified, and the second involves emulation of purses that have been properly certified. We consider the consequences of the second type of break-in, and determine how the exploitation of such a break-in is affected by system monitoring processes, accountability rules, and local processing rules. We further consider the interaction between transaction privacy and accountability.

Auditable Metering with Lightweight Security

Dahlia Malkhi, AT&T Labs - Research
dalia@research.att.com, http://www.research.att.com/~dalia

In this work we suggest a new mechanism for metering the popularity of web-sites: The compact metering scheme. Our approach does not rely on client authentication or on a third party. Instead, we suggest the notion of a_timing scheme_, a computation that can be performed incrementally, whose output is compact, and whose result can be used to efficiently verify the effort spent with high degree of confidence. We use the difficulty of computing a timing scheme to leverage the security of a metering method by involving each client in computing the timing function (for some given input) upon visiting a web site, and recording the result of the computation along with the record of the visit. Thus, to forge client visits requires a known investment of computational resources, which grows proportionally to the amount of fraud, and is infeasible for visit counts commonly found in the World Wide Web. The incremental nature of the timing function is used to create a new measure of client accesses, namely their duration.

Joint work with Matthew K. Franklin The URL for the published paper is: http://www.research.att.com/~dalia/pubs/fc97-ftp.ps.gz

Auditable Metering with Lightweight Security, by Dahlia Malki and Matthew Franklin. Financial Cryptography '97, (LNCS, 1318), R. Hirschfeld (Ed.), Anguilla, February 1997, pp. 151-160.

The Imperfection of Secrecy in Real Network Protocols

Hilarie Orman, DARPA
horman@darpa.mil

Authentication and key exchange protocols frequently rely on secret keys and randomly chosen information. The designs of such protocols often assume that the randomness is perfect and the keys are used for no other purpose. In practice, both assumptions are violated, and the strength of the protocol is reduced below its theoretical maximum. The talk will illustrate the problem and a proposed analysis method for deriving realistic estimates of the resistance of the protocols to attacks.

A Transparent Distributed Cryptographic Filesystem

Pino Persiano, University of Salerno
giuper@dia.unisa.it

In this talk I will present a cryptographic implementation of a Network File System which allows the use of private data stored on a server as a local resource without the need to trust the remote system..

This work is based on the following considerations:

1.The wide use of the client/server model, including the promising area of mobile computing. 2.The increasing need for privacy in data stored in a remote server. 3.The most used protocol, NFS, considered a standard, is really weak and suffers from many drawbacks which will be presented in the body of the paper. 4.All the previous implementations require a heavy user interface, which often make the system difficult to use for practical purposes.

Our filesystem, Transparent Cryptographic File System, runs on many Unix dialects providing the same user interface of SUN's NFS with DES, RC5 a,d IDEA encryption of all the data sent over the network. The full package is PD available (see `http://tcfs.dia.unisa.it`) and has been extensively tested in order to understand the performance impact due to the use of cypher and the robustness of the global system in a real size environment. I will also discuss the performance of the system and various directions for future work.

Efficient and Secure Metering

Benny Pinkas, The Weizmann Institute
bennyp@wisdom.weizmann.ac.il

We consider an environment in which many servers serve an even larger number of clients (e.g. the web), and it is required to meter the interaction between servers and clients. In particular, it is required to count the number of clients that approached a server. Possible applications are measurement of the popularity of web pages in order to decide on advertisement fees, and reliable usage based accounting between computer networks. The metering process must be very efficient for all the involved parties: clients, servers and the metering agency. It should also be secure against fraud attempts by servers who try to claim they served more clients than they actually did, or by clients who do not wish to help servers count their visits.

We suggest constructions for several metering systems, based on efficient cryptographic techniques. The proposed metering systems are efficient and secure. They are also very accurate and can preserve the privacy of clients.

Joint work with Moni Naor.

Hamiltonian Circuits on the n-Dimensional Cube

Henry Pollak, Teachers College of Columbia University
0006182700@mcimail.com

The problem of designing a counting circuit inside a computer turns out to require enumerating Hamiltonian circuits on the n-dimensional cube. We will talk about the history of the problem, see the solution in 3 and 4 dimensions, and get some idea why the problem is so tough for larger n.

Hacking and Networked Terrorism

Marcus Ranum, Network Flight Recorder, Inc.
mjr@clark.net, http://www.clark.net/pub/mjr

Society is increasingly accommodating networks into daily life, and with the advent of ubiquitous Internet access, hacking risks are becoming a new category of social problem. Today's law enforcement methods are generally incapable of making a dent in the hacking problem – will it get worse tomorrow? The speaker wishes to use this opportunity to explore some of the unpleasant parallels between hacking and terrorism, and the difficulty of solving social problems using software.

Problems with the Firewall Concept

Marcus Ranum, Network Flight Recorder, Inc.
mjr@clark.net, http://www.clark.net/pub/mjr

Internet firewalls are widely deployed as a front-line defense against intruders and hackers. First deployed widely in 1991, the firewall concept has evolved into a "must have" technology for many sites connecting to the Internet. Today, However, new technologies may be rendering the firewall obsolete. We will discuss some of the challenges that firewall technologies face, and attempt to extrapolate events in firewall technology over the next few years.

Toward Acceptable Metrics of Authentication

Mike Reiter, AT&T Labs - Research
reiter@research.att.com, http://www.research.att.com/~reiter/

Authentication using a path of trusted intermediaries, each able to authenticate the next in the path, is a well-known technique for authenticating entities in a large-scale system. Recent work has extended this technique to include multiple paths in an effort to bolster authentication, but the success of this approach may be unclear in the face of intersecting paths, ambiguities in the meaning of certificates, and interdependencies in the use of different keys. Several authors have thus proposed metrics to evaluate the confidence afforded by a set of paths. In this talk we present a set of guiding principles for the design of such metrics. We motivate our principles by showing how previous approaches fail with respect to them and what the consequences to authentication might be. We then propose a direction for constructing metrics that come closer to meeting our principles and thus, we believe, to being satisfactory metrics for authentication.

A paper on this work has been published as

M. K. Reiter and S. G. Stubblebine. Toward acceptable metrics of authentication. In Proceedings of the 1997 IEEE Symposium on Security and Privacy, pp. 10-20, May 1997.

The "Standards" Approach to K-12 Mathematics Education

Joseph Rosenstein, DIMACS - Rutgers University
joer@dimacs.rutgers.edu,
http://dimacs.rutgers.edu/nj_math_coalition/joer/joer.html

This session will include discussions of what standards are and why they have been developed by the National Council of Teachers of Mathematics and adapted in many states, the implications of the mathematics standards for the K-12 classroom, the potential benefits of the standards approach, and the concerns that have been raised about the standards. (High school teachers will be asked to share their experience with the standards.) New Jersey's Mathematics Standards and the NJ Mathematics Curriculum Framework (edited and co-authored by the presenter) will be highlighted as an example and model of mathematics standards and particularly, for its' treatment of discrete mathematics.

An article based on this presentation appears, with other articles, in
http://dimacs.rutgers.edu/nj_math_coalition/articles.joer/

Crowds: Anonymous Web Transactions

Aviel Rubin, AT&T Labs - Research
rubin@research.att.com, http://www.cs.nyu.edu/~rubin/

In this talk we introduce a system called Crowds for protecting users' anonymity on the world-wide-web. Crowds, named for the notion of "blending into a crowd", operates by grouping users into a large and geographically diverse group (crowd) that collectively issues requests on behalf of its members. Web servers are unable to learn the true source of a request because it is equally likely to have originated from any member of the crowd, and indeed collaborating crowd members cannot distinguish the originator of a request from a member who is merely forwarding the request on behalf of another. We describe the design, implementation, security, performance, and scalability of our system. our security analysis introduces degrees of anonymity as an important tool for describing and proving anonymity properties.

Joint work with Mike Reiter.

`http://www.research.att.com/projects/crowds`

How Math Will Protect You on the Information Super-highway

Aviel Rubin, AT&T Labs - Research
`rubin@research.att.com`, `http://www.cs.nyu.edu/~rubin/`

In this talk, I'll discuss how some simple mathematical techniques can be used to achieve strong cryptography. Several examples will be used.

Mobile Code Security Issues

Fred B. Schneider, Cornell University
`fbs@cs.cornell.edu`
`http://www.cs.cornell.edu/Info/Department/Annual96/Faculty/Schneider.html`

When programs (known as agents) can roam a network of hosts, three security concerns must be addressed:

i.host integrity – ensuring that execution of an agent cannot compromise the hosts on which that agent executes

ii.agent integrity – ensuring that a computation structured from agents will be correctly completed, despite the existence of faulty or malacious hosts in the network

iii.inter-agent security – ensuring that agents are not compromised in their interactions with other agents.

Solutions to these problems are being investigated in connection with the TACOMA Too System, an ML-based environment for constructing systems of agents. This talk will outline the solutions being explored. The work on host integrity involves novel applications of cryptographic abstractions to implementing fault-tolerance. The scheme for specifying inter-agent security policies involves a new and expressive mechanism.

The paper has since been published:

Schneider, F.B., "Towards fault-tolerant and secure agentry". Invited paper. Proc. 11th International Workshop WDAG '97 (Saarbrucken, Germany, Sept. 1997), Lecture Notes in Computer Science, Volume 1320, Springer-Verlag, Heidelberg, 1997, 1-14.

Copies of the paper are also available through www:
`http://www.cs.uit.no/DOS/Tacoma/Publications.html`

Cryptanalytic Fault Attacks

Adi Shamir, The Weizmann Institute
`shamir@wisdom.weizmann.ac.il`

In this talk I'll describe new methods for extracting cryptographic keys from sealed tamper-resistant devices such as smart cards, by inducing either temporary or permanent faults into their cryptographic computations.

A Formal Treatment of Transactional Trust Management

Martin Strauss, AT&T Labs - Research
`mstrauss@research.att.com`, `http://www.research.att.com/~mstrauss`

A trust management engine processes requests (like "purchase a computer"), supporting credentials (like "Vice President Smith says 'Alice is authorized to purchase a computer' "), and policies (like "all purchases must be approved by a vice president.") In general, credentials and policies are fully-expressive programs. The job of the trust management engine is to decide whether or not the supporting credentials prove that the request complies with the policy and, in some cases, to say why a non-compliant request fails to comply. Trust management is a crucial component of many network services, particularly those that use public-key cryptography to process signed requests. The notion of trust management was introduced in [BFL], as was the PolicyMaker trust management engine.

In this paper, we present a formal model that captures the notion of "trust management" embodied in PolicyMaker. We investigate this notion of "transactional" trust management from a complexity-theoretic point of view, showing, for example, that the question of whether a set of credentials prove that a request complies with policy is undecidable in general, and NP-hard even if restricted in any of a number of straightforward ways. We give necessary and sufficient conditions on the input so that the question is solvable in polynomial time and has some other desirable features. Finally, we point out some limitations of the "transactional" model of trust management and mention some on-going work on alternative designs for general-purpose trust management engines.

This is joint work with Joan Feigenbaum.

[BFL] M. Blaze, J. Feigenbaum, and J. Lacy. "Decentralized trust management", in Proceedings of the IEEE Symposium on Security and Privacy, pages 164–173, Oakland CA, May 1996.

Reliable and Private Communication Over Echo Lines

Rebecca Wright, AT&T Labs - Research
`rwright@research.att.com, http://www.research.att.com/info/rwright`

We present and explore the echo communication model. Nodes of a graph communicate via echo, in which the same message is sent to all their neighbors. An adversary has control of some subset of the nodes, considered "faulty". Both correct and faulty nodes are constrained by the echo property. We present efficient protocols for almost perfectly reliable message transmission and perfectly secret message transmission. In addition, we give a perfectly secret, but inefficient, message transmission and show that perfectly reliable message transmission is impossible.

Joint work with Matthew Franklin.

Overview of Secure Co-Processors

Bennet Yee, University of California at San Diego
`bsy@cs.ucsd.edu, http://www-cse.ucsd.edu/~bsy/`

Secure coprocessors are suicidal devices, microcomputers with a death-wish. Designers of these suicidal machines make them very sensitive to certain special classes of external events, triggering self-destruction. These coprocessors are extra hardware modules that may be plugged into existing workstations or PCs, and provide the machines with the following special property: even though many people, including those who are not entirely trustworthy, may have physical access to the machines, security properties derived from the secure coprocessor will not be compromised.

In this talk, I will give an overview of the packaging technology used to protect secure coprocessors, discuss what security properties may be "bootstrapped" from a secure coprocessor to the system, go over some applications of secure coprocessors, and talk about some of the limitations of this technology.

Joint work with Doug Tygar.

The Sanctuary Project

Bennet Yee, University of California at San Diego
`bsy@cs.ucsd.edu, http://www-cse.ucsd.edu/~bsy/`

In the research community, there has been some interest in the idea of mobile agents: software that autonomously migrate from one server machine to the next, performing services on their owner's behalf. These services range from finding the best price on some commodity (e.g., airline tickets, music CDs) to automatically (and intelligently) summarizing web page contents and news articles for their owners.

There are significant benefits to this computation model: while computers and networks only get faster with improved technology, the speed of light won't. This means that the traditional client-server model, where clients would make repeated accesses to distant resources (e.g., web servers, database servers), is inherently performance limited. If, instead of remotely querying for data, we could move the code (agent) to the data (server), we could gain tremendously in performance. (The code size is typically much smaller than the data size.)

Unfortunately, there are also many security problems that arise with the mobile code model: not only might an untrustworthy agent violate the servers' security, but an untrustworthy server might also modify the result of computation done at a previous server – and thus falsify the returned result.

In this talk, I will give an overview of the Sanctuary project, discuss several practical approaches to protecting the results computed at remote servers, and discuss some related problems.

Papers' URLs are:

`http://www-cse.ucsd.edu/~bsy/pub/th.ps.gz.sanctuary.ps`