Research in Cryptography, Information Security, and Algorithm Development
9807-12&26

Progress Report: July 1, 2001—December 31, 2001

Shafi Goldwasser, Ronald L. Rivest and Michael Sipser

Project Overview

Recent NTT-sponsored research has focussed on both designing secure cryptographic protocols and showing their optimality by exhibiting impossibility of stronger cryptographic protocols in both multi-party (distributed) and two-party scenarios.

The emphasis of the research, which is expressed in each one of our projects, is on the development of novel frameworks and theoretical provably secure solutions to problems arising from applications.

Progress Through December 2001

A. SECURE MULTI-PARTY PROTOCOLS WITH MANY FAULTY PROCESSORS

Professor Shafi Goldwasser, jointly with Yehuda Lindell from the Weizmann Institute of Science, worked on secure multi-party protocols with many faulty processors.

The beautiful result of Lindell, Lysyanskaya, and Rabin which showed that authenticated Byzantine agreement does not compose (concurrently or in parallel) when the number of faults exceeds half of the processors, puts in question the entire family of results on secure multi party protocols when the number of faults exceeds half of the processors, as these results assume broadcast (which is implemented by running Byzantine agreement) as a building block. In new work with Lindell we show how to relax the definition of Byzantine agreement so that the relaxation can be achieved with more than a half faults. The relaxation requires that all honest parties will either receive a special "Faulty protocol" output, or the "correct output" but never will there be a disagreement between honest parties as to the value of the correct output. Extending this to any secure multi party computation, we show how to transform any protocol which is secure assuming conventional broadcast into one that is secure with the new broadcast and in which each honest processor will either receive the correct output or output "faulty protocol'. The new protocols do compose concurrently.

B. CRYPTOGRAPHY BASED ON LATTICE PROBLEMS

Together with Professor Daniele Micciancio from University of California at San Diego, Professor Shafi Goldwasser has written a book titled "Complexity of Lattice Problems: A Cryptographic Perspective" in which all known results in lattice based cryptography and cryptanalysis are presented. We believe this is a fundamentally important area that has not received proper treatment from a theory of computation perspective before.

C. ROBUST MIX-NETS FOR ELECTRONIC VOTING

Professor Ron Rivest, together with Ari Juels and Markus Jakobsson of RSA Security, have proposed a new technique for making mix nets robust, called randomized partial checking (RPC). The basic idea is that rather than providing a proof of completely correct operation, each server provides strong evidence of its correct operation by revealing a pseudo-randomly selected subset of its input/output relations.

Randomized partial checking is exceptionally efficient compared to previous proposals for providing robustness; the evidence provided at each layer is shorter than the output of that layer, and producing the evidence is easier than doing the mixing. It works with mix nets based on any encryption scheme (i.e., on public-key alone, and on hybrid schemes using public-key/symmetric-key combinations). It also works both with Chaumian mix nets where the messages are successively encrypted with each servers' key, and with mix nets based on a single public key with randomized re-encryption at each layer.

Randomized partial checking is particularly well suited for voting systems, as it ensures voter privacy and provides assurance of correct operation. Voter privacy is ensured (either probabilistically or cryptographically) with appropriate design and parameter selection. Unlike previous work, our work provides voter privacy as a global property of the mix net rather than as a property ensured by a single honest server. RPC-based mix nets also provide very high assurance of a correct election result, since a corrupt server is very likely to be caught if it attempts to tamper with even a couple of ballots.

We note that our proposal might be usable to implement the "anonymous channels" required by some recent voting scheme proposals, such as the voting scheme described by Tatsuaki Okamoto in his paper, ``Receipt-Free Electronic Voting Schemes for Large Scale Elections.''

[JJR02] ``Making Mix Nets Robust for Electronic Voting by Randomized Partial Checking.'' By Ari Juels, Markus Jakobsson, and Ronald L. Rivest. (submitted)

D. TWEAKABLE BLOCK CIPHERS

Professor Ron Rivest, together with Moses Liskov, have proposed a new cryptographic primitive---the ``tweakable block cipher.'' Such a cipher has not only the usual inputs---message and cryptographic key---but also a third input, the ``tweak.'' The tweak serves much the same purpose that an initialization vector does for CBC mode or that a nonce does for OCB mode. Our proposal thus brings this feature down to the primitive block-cipher level, instead of incorporating it only at the higher modes-of-operation levels. We suggest that (1) tweakable block ciphers are easy to design, (2) the extra cost of making a block cipher ``tweakable'' is small, and (3) it is easier to design and prove modes of operation based on tweakable block ciphers. Moreover, we can prove tighter security

bounds for certain modes of operation based on tweakable block ciphers than are known for the corresponding modes based on standard block ciphers (e.g. for OCB mode).

[LR02] ``Tweakable Block Ciphers.'' By Moses Liskov and Ronald L. Rivest. Submitted.

E. AUTHENTICATED BYZANTINE AGREEMENT

Work by Anna Lysyanskaya, together with Yehuda Lindell of the Weizmann Institute of Science, and Tal Rabin of IBM T.J.Watson Research; have worked on authenticated Byzantine agreement.

A classical result is that in order to reach agreement in a point-to-point network with t malicious participants, the total number *n* of players must be *n > 3t+1*. It is also well known that augmenting the network with a public-key infrastructure results in protocols for any *n > t*. This augmented problem is called ``authenticated Byzantine agreement.''

Here, we consider whether several copies of authenticated Byzantine agreement can be executed when *n < 3t+1*. We discover that concurrent execution of multiple copies of the protocol is impossible. In the case of sequential execution, we discover that it is impossible for deterministic protocols with a bounded number of rounds. In contrast, randomized protocols that compose sequentially are possible. We exhibit two such protocols: one for the case when *2t < n < 3t+1*; the other for the case *n > t*, *n = O(log k)*, where *k* is the security parameter. Finally, we show that any number of authenticated Byzantine agreement protocols for any *n > t* can be concurrently invoked if each of them is given a unique identifier.

[LLR02] "On the Composition of Authenticated Byzantine Agreement." By Yehuda Lindell, Anna Lysyanskaya, and Tal Rabin. To appear in STOC'02.

F. SEQUENTIAL COMPOSITION OF PROTOCOLS WITHOUT SIMULTANEOUS TERMINATION

Anna Lysyanskaya, jointly with Yehuda Lindell of the Weizmann Institute of Science, and Tal Rabin, IBM T.J.Watson, considered the question of sequential composition of protocols without simultaneous termination.

The question of the composition of protocols is an important and heavily researched one. We consider the problem of sequential composition of protocols that do not have simultaneous termination, i.e. parties may complete the execution at different times. An important example of a protocol that has this property is that of randomized Byzantine Agreement (with an expected constant number of rounds). Given that the termination of the parties is not simultaneous, a natural question to consider is how to synchronize the parties so that such protocols can be sequentially composed. Furthermore, such a composition should preserve the original running time of the protocol, i.e. running the protocol *L* times sequentially should take in the order of *L* times the running time of the protocol.

In this paper, we present a method for sequentially composing any protocol in which the players do not terminate in the same round, while preserving the original running time. An important application of this result is the

sequential composition of parallel Byzantine Agreement. Such a composition can be used by parties connected in a point-to-point network to run protocols designed for the broadcast model, while maintaining the original round complexity.

[LLR02a] "Sequential Composition of Protocols without Simultaneous Termination." By Yehuda Lindell, Anna Lysyanskaya, and Tal Rabin. Submitted.

G. UNIQUE SIGNATURES AND VERIFIABLE RANDOM FUNCTIONS FROM THE DH-DDH SEPARATION.

Anna Lysyanskaya has come up with a new construction of a unique signature scheme.

Unique signature schemes, also known as invariant signature schemes, were introduced by Goldwasser and Ostrovsky. These are signature schemes where a signature is a (hard) function of pk and m, for all, even adversarially chosen, pk. The only previously known construction of unique signatures in the plain model was based on the RSA assumption. Constructions in the common-random-string model and in the random-oracle model also existed.

We give a simple construction of unique signatures based on an assumption about groups in which decisional Diffie-Hellman is easy, and computational Diffie-Hellman is hard. Several recent results suggest plausibility of such groups.

It is important to note that our signature scheme is simple, natural, and stateless, and therefore is a step forward not only as far as unique signatures are concerned, but also as far as practical signature schemes are concerned. To date, the only practical stateless provably secure signatures (other than in the random-oracle model) were based on the strong RSA assumption; our signature scheme is based on a different assumption.

An important implication of this result is the new construction of verifiable random functions (VRFs). VRFs, introduced by Micali, Rabin, Vadhan are objects that combine the properties of pseudorandom functions (i.e. indistinguishability from random even after querying) with the verifiability property. These objects are very handy for cryptographic protocol design, because they can be viewed as a commitment to an arbitrary number of bits. Yet, until this work, only one candidate construction for this primitive existed.

[L02] “Unique Signatures and Verifiable Random Functions from the DH-DDH Separation.” By Anna Lysyanskaya. Submitted.

H. TAMPER-PROOF SECURITY

Anna Lysyanskaya and Professor Silvio Micali, jointly with Tal Malkin of AT&T Research and Rosario Gennaro and Tal Rabin of IBM T.J.Watson, worked on tamper-proof security.

A cryptographic algorithm is executed by various parties, where one (or more) has some secret information. The classical security definitions for these algorithms assume that an adversary has no access whatsoever to the secret information of honest parties. Rather, the adversary is only allowed to query the cryptographic algorithms on inputs of its choice, where the answer is always computed according to the correct original secret information.

However, many of today's cryptographic applications are carried out on small portable devices, such as smartcards. These devices try to offer physical security features that prevent the adversary from reading the secrets stored inside the card. However, as recent research has shown, it is possible for the adversary to partially manipulate the secret key by injecting faults into the storage. This might have catastrophic effects on the security of the algorithms implemented by the device.

In this paper we propose a new security model that deals with this type of attacks. Our model allows the adversary to apply a function *f* to the secret key *sk* and obtain the result of the cryptographic algorithms using the new secret key *f(sk)*.

For the most common cryptographic applications, such as signature schemes, cryptosystems, and zero knowledge proofs, we formally define tamper-proof security, i.e., security in this new model. We then show that for tamper-proof security against general polynomial-time functions *f*, it is necessary to enhance the device with a self-destruct capability, as well as with public parameters. In this enhanced model, we give constructions of tamper-proof signatures and cryptosystems with respect to general functions *f*. We also show that with respect to general functions *f*, it is impossible to construct tamper-proof zero-knowledge proofs even in the enhanced model.

Finally, we consider several interesting specific functions and achieve trapper-proof security with respect to them. This latter results are interesting in that for some specific and yet powerful functions, we can achieve tamper-proof security for all three cryptographic applications considered, and even without the some of the enhancements to the model.

[GLMMR02] “Tamper-Proof Security.” By Rosario Gennaro, Anna Lysyanskaya, Tal Malkin, Silvio Micali, and Tal Rabin. Submitted.

I. PROPERTY-TESTING

Sofya Raskhodnikova has worked in the area of property testing.

The field of property testing studies algorithms that distinguish, using a small number of queries, between inputs which satisfy a given property, and those that are `far' from satisfying the property. Testing properties that are defined in terms of monotonicity has been extensively investigated, primarily in the context of the monotonicity of a sequence of integers, or the monotonicity of a function over the *n*-dimensional hypercube {1..m}^n. These works resulted in monotonicity testers whose query complexity is at most polylogarithmic in the size of the domain.

We show that in its most general setting, testing that Boolean functions are close to monotone is equivalent, with respect to the number of required queries, to several other testing problems in logic and graph theory. These problems include: testing that a Boolean assignment of variables is close to an assignment that satisfies a specific

2-CNF formula, testing that a set of vertices is close to one that is a vertex cover of a specific graph, and testing that a set of vertices is close to a clique.

We then investigate the query complexity of monotonicity testing of both Boolean and integer functions over general partial orders. In particular, we show that there exist partial orders for which non-adaptive monotonicity testing cannot be performed with a polylogarithmic number of queries; in the course of this proof we construct graphs with combinatorial properties that may be of independent interest. Because of the aforementioned relationships, this implies the same lower bounds on several other testing problems. We also show that in the special case of testing functions over the hypercube, the number of required queries has to depend on the dimension.

On the positive side, we show that even in the most general setting, the number of required queries is upper bounded by the square root of the input size. Furthermore, we provide testing algorithms with significantly improved efficiency for several subclasses of partial orders whose Hasse diagrams satisfy various connectivity properties.

J. CRYPTOGRAPHY OVER FIELDS WITH UNBOUNDED COMPUTATION

David Woodruff and Marten van Dijk wrote a paper "Cryptography in an Unbounded Computational Model" for Eurocrypt 2002. This paper shows that it is not possible to share a secret in the proposed model of computation (where infinite-precision field operations are assumed to have unit cost, and where each party can generate random real numbers). This immediately rules out many cryptographic primitives, such as Oblivious Transfer, Diffie-Hellman Key Exchange, interactive encryption, etc. These are stunning results for this new model of computation.

For more details, see section 6 of the final paper at web.mit.edu/~dpwood/www/euro.ps. The proof uses a lot of field theory. The idea is to show that after each atomic step of the protocol, we have as an invariant that the intersection of Alice's field with Bob's field equals the field known by the public (and eavesdropper). The two atomic steps are: (1) a party samples a random real number and (2) a party transmits an element of his/her field to the other party. After each step a degree-based invariant is maintained which in turn implies the invariant.

[WvD02] “Cryptography in an Unbounded Computational Model.” By David Woodruff and Marten van Dijk. To appear in EUROCRYPT’02.

K. FINGERPRINTING

Christopher Peikert, Abhi Shelat, and Adam Smith worked on fingerprinting.

Fingerprinting is a technique for protecting published data against unauthorized copying by making each copy unique. For example, early vendors of logarithm tables introduced unique, unnoticeable errors in the each copy of the tables. If a buyer ever sold illegal copies, the true owner could trace the source and have him prosecuted.

Today, creators of digital content are attempting to protect their works in a similar way. Each copy of a work is "marked" in some way that makes no perceptual difference (e.g., without changing the appearance of an image or

the sound of a song), yet which identifies it uniquely. Of course, uniqueness is not sufficient: the creator must also make these marks difficult to detect, and even guard against a coalition of users who compare their copies to discover the marks. In order to abstract away the details of embedding these changes, we consider the following problem: design a collusion-secure code so that if a small number of users combine their codewords into one unauthorized copy, it can be traced back to at least one of the guilty parties.

Several definitions of collusion-secure codes have been proposed, each with varying strengths. Most definitions allow the code's alphabet to be arbitrarily large, but because we are concerned with protecting digital data, we consider the model proposed by Boneh and Shaw (Crypto 1991). In it, the alphabet of the code is binary, the coalition of users can only detect and modify positions for which their codewords differ, and the distributor's tracing algorithm must identify one of the guilty users with high probability (taken over the random choices of both the distributor and the coalition).

One of the most important parameters is the maximum allowable size, *c*, of a coalition. Boneh and Shaw give a construction of a collusion-secure code whose length is cubic in *c*, but the total number of codewords is only linear in the length. By concatenation, they achieve a code with exponentially-many codewords whose length is quartic in *c*. Conversely, they prove that any collusion-secure code must have length at least linear in *c*.

It is not unreasonable to imagine that dozens (if not hundreds) of users might collude to defeat a fingerprinting scheme, and in this scenario the Boneh-Shaw construction may be too inefficient. We are interested in narrowing the gap between the construction and lower bound, either by demonstrating more efficient codes or showing that no such objects exist. To improve the lower bound requires a general strategy, to be employed by a coalition, for defeating any (short enough) fingerprinting code.

In the Boneh-Shaw model, when *c* is very large, the risk to each individual in a coalition is relatively low. We are therefore also interested in studying codes that allow the distributor to trace more than one (e.g., a constant fraction) of the guilty parties with high probability. Because it is possible for only one member of the coalition to contribute to the copy, we must modify the model and definitions to make this investigation meaningful. We have begun this work, and feel we are making good progress.

L. CRYPTOGRAPHIC PROTOCOL COMPILER

Stephen Weis and Ioannis Tsoukalidis have recently developed a new tool for implementing cryptographic protocols. This tool allows one to enter a protocol in a special-purpose high-level "cryptographic protocol language", and then to "compile" this language to produce (1) running Java code; (2) Latex protocol diagrams. This provides a "fast track" for developing prototype implementations of newly designed cryptographic protocols, and for providing documentation. This work is done in conjunction with Professor Rinard's compiler group.

This work is documented further at: http://theory.lcs.mit.edu/~cis/cpl/index.html .

M. CRYPTOGRAPHY WITH QUANTUM DATA

Group Members: Adam Smith

Quantum computers, if they could be built at practical scales, would be able to quickly solve problems that are not known to be solvable on a traditional computer. This would affect classical cryptography, since most known public-key systems would become easily breakable. It would also introduce a variety of completely new challenges, since quantum computers represent data very differently from classical computers. This project focuses on understanding what cryptographic tasks can be performed with quantum data as inputs.

Research has focused principally on three tasks:

- Data authentication [BCGST]:

  Guaranteeing authenticity of transmissions over insecure channels is important in distributed applications. However, traditional techniques do not suffice to authenticate the transmission of quantum states. We develop an efficient, information-theoretically secure quantum authentication scheme (QAS). We show that unlike in the classical case, any QAS must also encrypt the data it authenticates. This has several consequences. On one hand, it shows that our scheme is essentially optimal in both message size and key length. It also shows that digital signatures of quantum states are impossible, even if we only require security against computationally limited adversaries.

  Our construction introduces the notion of a purity-testing protocol, which is of independent interest and is also used for cryptographic entanglement purification.

- Distributed Computing [CGS]:

  Multi-party computing, also called "secure function evaluation", has been extensively studied in classical cryptography. We consider the following general extension of this task: players each provide a quantum state as input, and collectively evaluate a quantum circuit in such a way that a cheating coalition can do no more than choose their inputs and receive their outputs.

  Our protocols are information-theoretically secure, i.e. no assumptions are made on the computational power of the cheating coalition. For the slightly weaker task of "verifiable quantum secret sharing", we give a protocol which tolerates any $t < n/4$ cheating parties (out of n). This is shown to be optimal. We use this new tool to show how to perform any multi-party quantum computation as long as the number of dishonest players is less than $n/6$.

- Entanglement Purification [ASY]:

  Many quantum cryptographic protocols, including those for key distribution and authentication, implicitly contain an "entanglement purification" phase in which good Einstein-Podolsky-Rosen (EPR) pairs are established between two participants. We give a precise definition for this task as it is needed in these settings. We show that error-free purification protocols cannot achieve good parameters. For protocols with small error, we show that the purity-testing protocols developed for authentication [BCGST] are essentially optimal.

[BCGST02] H. Barnum, C. Crépeau, D. Gottesman, A. Smith and A. Tapp. "Authentication of Quantum Messages". Manuscript.

[CGS02] C. Crépeau, D. Gottesman and A. Smith. "Quantum Multi-party Computation". To appear in Proceedings of STOC 2002, ACM, May 2002.

[ASY02] A. Ambainis, A. Smith and K. Yang. "General Entanglement Purification Protocols". To appear in Proceedings of CCC 2002, IEEE, May 2002.

Research Plan for the Next Six Months

Over the next six months, we plan to continue our work on the both the fundamentals and the applications of cryptography to information security. We tend to work on a wide variety of problems simultaneously; it is hard to predict exactly which problems and approaches are likely to be fruitful. However the following directions will certainly be emphasized over the next six months.

- National ID card schemes: We propose to investigate national ID card schemes, to see how individual privacy can be maintained while providing some of the benefits that such cards might offer. These cards are rather controversial in the U.S., and further research might help point the way to technical solutions to some of the vexing issues. Privacy in general is one driving motivation for our research program.

- Proof techniques for cryptographic protocols: We continue to explore expanded models of adversarial capability, to design new cryptographic protocols to defeat such adversaries, and to invent new proof techniques for proving such protocols to be secure. Some of this work may go in the direction of the recent Abadi-Rogaway proposal to unify classic cryptographic proof techniques with more standard program-proving techniques from logic; issues such as using keys to encrypt themselves are particularly hard to handle.

- User authentication in pervasive computing environments: This problem is particularly difficult when the basic components (e.g. cell phones or laptops) may be anonymous commodity-like devices that the user picks up or borrows as needed. We are studying schemes for user authentication involving user-worn beacons that broadcast cryptographic information to these insecure devices that can then act on the user's behalf. We have made some progress on such schemes, but much remains to be done.

- Fundamental cryptographic primitives: There remains a rich store of open problems in the area of fundamental cryptographic primitives. Recent advances, e.g. in all-or-nothing transformations and block ciphers, suggest many open problems. For example, there are generalizations of the Luby-Rackoff paradigm that remain intriguing and unsolved: when does a generalized network of Feistel-like transformations yield a pseudo-random permeation or a super pseudo-random permutation? Our recent work on "tweakable block ciphers" (see above) gives rise to a multitude of questions of both practical and theoretical importance. Further extensions to our work on cryptography over fields may yield greater insight into the structure of cryptographic systems, perhaps yielding new approaches for achieving security in the real world.

- Security of routing protocols: We continue to be concerned about the security of the infrastructure of the Internet, and are studying ways to provide much-needed authentication to the underlying routing protocols. (In particular, the problem of properly authenticating the routing tables if of interest to us.) As it stands now, the Internet is much too fragile and vulnerable to attack at its core routing level; we would like to provide new defenses against malicious routers who attempt to disrupt or bring down the Internet by distributing spurious information.