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

Shafi Goldwasser, Ronald L. Rivest and Michael Sipser

Project Overview:

We propose cutting-edge research in security and quantum computation in areas we believe to be of mutual interest to researchers in LCS and NTT. While some of this represents a continuation of research previously supported by NTT, much of it is new. This is a brief summary of our research plan for July 1, 2000 through June 30, 2001.

It is perhaps important to mention at this point that while the concrete research directions proposed below are firm ones, we may also wish to use NTT support to explore related directions not mentioned here. In particular, we note that there are likely to be several new graduate students in our group next year, and their precise research directions are as usual subject to negotiation and creativity.

I. Threshold Cryptography

We will continue research in threshold cryptography—cryptography that enables a set of dedicated servers to simulate a trusted entity, assuming that at most a fraction of a set of dedicated servers have been corrupted. Such a trusted entity can serve as a trusted third party in protocols, for example, in certification of public keys, in fair exchange, etc. This research is important both in theory and in practice; it is the cryptographic equivalent of achieving high reliability in systems through the systematic use of redundancy (except here the goal is security rather than reliability).

We have developed a general technique for guaranteeing security of threshold schemes against the adversary that can adaptively corrupt a number of servers, and which also guarantees security when many instances of the protocol are run concurrently [1,3]. We have also developed specific threshold constructions for several important cryptographic primitives (RSA and ElGamal cryptosystems, DSS, Cramer-Shoup cryptosystem) without having to assume ``secure erasure" [1,2,4].

We plan to work on establishing lower bounds on communication in threshold systems; basing efficient threshold cryptography on general assumptions; and investigating the possibility of threshold cryptography in more general adversarial models.

A resource which is heavily used in [4,3] is for the dedicated servers (simulating the trusted entity) to share in advance— in an off-line set-up stage — randomness to be used later during the on-line protocol. Unfortunately, even though the ability to share randomness in advance off-line enables proving strong cryptographic security properties, it does require secure storage of quite a bit of data. We plan to work on eliminating the need of shared randomness in a set-up stage and hopefully replacing it by pseduo-random generation on-line.

We will also work on important related open problems in general multi-party computation, such as secure multi-party computation that remains secure under concurrent composition (this has applications in secure electronic transactions) and more efficient erasure-free multi-party computation. Some of our threshold techniques described above may be applicable here. Finally, we plan to explore receipt-free multi-party computation, which has particular application to electronic voting.

For more information, see http://theory.lcs.mit.edu/~cis/cis-threshold.html

[1] Stanislaw Jarecki, Anna Lysyanskaya. "Adaptively secure threshold cryptography: introducing concurrency, removing erasure." In Advances in Cryptology — Proceedings of Eurocrypt2000, Springer-Verlag, to appear, May 2000.

[2] Stanislaw Jarecki, Anna Lysyanskaya "Adaptively secure threshold cryptography without erasures." Manuscript, 2000.

[3] Anna Lysyanskaya. "Threshold cryptography secure against the adaptive adversary, concurrently." Manuscript, 2000.

[4] Ran Canetti, Shafi Goldwasser. "An efficient threshold public-key cryptosystem secure against adaptive chosen ciphertext attack." Appeared in EUROCRYPT'99, pp.90-106

II. Problems in the Theory of Block Cipher Design.

We also plan to consider efficient designs for provably secure block ciphers. Currently, there is a large gap between the efficiency of provably secure block ciphers and the efficiency required in practice, which leads to the use of unproven (albeit apparently secure) designs. It is therefore desirable to narrow this gap.

In practice, block cipher design is ad-hoc, and there is little mathematical evidence for security. Luby and Rackoff[1] were the first considered the problem of designing provably secure block ciphers. They succeeded, but their construction is impractical. In a number of recent breakthroughs [2,3,4] we have built upon their work.

We have not only made their construction practical, but we have also greatly deepened our understanding of such Luby-Rackoff ciphers. We now understand the importance of the types of underlying operations used, the cryptographic properties needed for security during some of the computations, and the significance of the ‘round functions’ used in these ciphers. These breakthroughs have given us valuable insight that will enable us to tackle a variety of new problems; our goal is to build a theory for block cipher design as a whole. We will proceed by exploring the use of weaker round functions in Luby-Rackoff ciphers and the use of alternate underlying designs.

[1] M. Luby and C. Rackoff. How to construct pseudorandom permutations from pseudorandom functions. SIAM J. Computing 17(2):373-386, April 1988.

[2] S. Patel, Z. Ramzan, and G. Sundaram. Towards Making Luby-Rackoff Ciphers Optimal and Practical. Proceedings of Fast Software Encryption 1999.

[3] S. Patel, Z. Ramzan, and G. Sundaram. Luby-Rackoff Ciphers: Why XOR is not so Exclusive.

[4] Z. Ramzan and L. Reyzin. On the Round Security of Symmetric Key Cryptographic Primitives.

III. Multi-Signer Signature Schemes

We plan to continue research on various paradigms for multi-signer signature schemes and their efficient implementations. Our most recent work (“Provable-Subgroup Signatures'', by Leo Reyzin (MIT), Silvio Micali(MIT), and Kazuo Ohta(NTT)) enabled any subgroup of signers to efficiently sign a message so that the signature provably reveals and commits the signing subgroup. This may have significant practical applications.

Because multi-signer signature schemes are often used in high-security applications, it is also natural to consider incorporating forward security into such schemes, where disclosure of secrets does not invalidate previous signatures. No forward-secure multi-signer schemes are known so far; we hope to achieve a breakthrough on this important problem.

IV. Secure Distributed Assertion Infrastructure

We have previously explored techniques for building a powerful and flexible public-key infrastructure. (This is the “SPKI/SDSI'' proposal described in http://theory.lcs.mit.edu/~cis/sdsi.html.) This work is maturing, and we have built a prototype secure web server based on SPKI/SDSI.

We now would like to explore more ambitious distributed assertion infrastructures that are more expressive, yet in which one retains the ability to reason and draw inferences securely. We have begun preliminary discussions with Tim Berners-Lee and colleagues on such a “semantic web'' proposal, and look forward to developing a design for a “semantic web'' in conjunction with them. A particular first step for such a design would be to check that a large organization could express and implement easily an access-control policy, such as for web access. Later steps might involve reasoning about rights to be given to code that has been downloaded.

(Note: some research on “semantic web'' has been concurrently proposed for support to DARPA, with Tim Berners-Lee as the PI.)

IV. Quantum Computation

Quantum computers may exploit quantum mechanical phenomena to achieve great speed when solving certain problems. If they could be built, a quantum computer could achieve the effect of great parallelism by using the phenomenon of superposition. Peter Shor and Lov Grover have shown that this parallelism might be quite useful for solving problems such as factoring integers and general searching.

In the past two years we (Professor Sipser, working together with Eddie Farhi and Jeffrey Goldstone of the MIT Physics department and Sam Gutmann of the Northeaster University Mathematics department) have obtained results that have helped to clarify the power of quantum computers. We have shown that quantum computer are essentially no faster than classical computers on certain parity and insertion problems. in contrast with searching problems that have been shown by Grover to admit a square-root speedup on quantum computers.

We are currently exploring ways of improving upon Grover's algorithm. Several years ago, Farhi and Gutmann gave a geometrical analysis of Grover's algorithm that is much simpler than Grover's original cumbersome analysis. Brassard, Hoyer, and Tapp extended Grover's algorithm to count the number of items satisfying a given condition in O(sqrt(N) steps. Kazuo Ohta of NTT Laboratories and Michael Sipser are working on a paper that gives a simpler Farhi-Gutmann type of analysis for the quantum counting algorithm.

Grover's algorithm has been shown to be optimal in the relativized sense. Thus, among algorithms that solve the searching problem using black-box queries, none can improve upon Grover's square-root speedup. In the past decade in complexity theory, a number of classical algorithms have been developed that avoid relativized limits. These algorithms operate outside the black-box paradigm, using algebraic methods to solve such problems as the counting satisfiability problem on an interactive proof system. These non-relativizing algorithms suggest the possibility that related methods may yield quantum algorithms that transcend the relativized limitations, and in so doing improve upon Grover's algorithm.

In a recent paper we have developed an algorithm for the satisfiability problem that relies upon adiabatic evolution to find a satisfying assignment. This method appears to use the power of quantum computation in a new way, and does not appear to relativize. Thus far we have been unable to analyze the running time of this algorithm. One of our goals for the coming period is to provide an analysis, as well as to investigate other avenues for speeding the search process through quantum computation.

Adam Klivans, a student in our group, has been considering quantum algorithms for problems in computational learning theory. One motivation for this research direction is to expand the range of examples for which quantum computers may outperform classical computers. He is working on the problem of whether the class of DNF formulae are learnable under the uniform distribution by a quantum computer given examples from a classical oracle. Conceivably, DNF formulae may be learnable under any distribution in the quantum PAC setting. The main technical difficulty lies in randomly sampling a given function's Fourier expansion with respect to a non-uniform distribution. We also believe this to be a useful technical question in its own right.