Research in Cryptography, Information Security and Algorithm Development
9807-12&26
Progress Report: July 1, 1999—December 31, 1999
Shafi Goldwasser, Ronald Rivest and Mike Sipser

1 Project Overview

Recent NTT-sponsored research has focussed on the topics of secure electronic voting, quantum computation, probabilistic property testing, threshold encryption and signature schemes, and derandomization techniques .

The emphasis of our reseach which is expressed in our work on each of the above topics is on the development of novel frameworks and theoretical provably secure solutions to problems arising from applications.

In this report, we will elaborate on our progress and plans in quantum computation, threshold cryptography, signature schemes and property testing.

1.1 Quantum Computation

Quantum computers are devices, which, in principle, can take advantage of quantum mechanical phenomena to achieve great speed when solving certain problems. If such computers could be built, they would use the phenomenon of superposition to achieve the effect of high parallelism within a single sequential processing engine. The recent theoretical work of Peter Shor and Lov Grover has demonstrated that this parallelism might be useful in solving problems of interest, such as factoring integers and general searching.

We are investigating the theoretical aspects of quantum computers along two avenues. First, we are attempting to devise algorithms to exploit the power of quantum computers in new ways to broaden the class of problems where a quantum speedup may be achieved. Second, we are studying the theoretical limitations of quantum computers through proofs of inherent computational difficulty on models of quantum computation.

1.2 Threshold Cryptography

The idea of theshold cryptography is to protect information (or computation) by fault-tolerantly distributing it among a cluster of cooperating computers. First consider the fundamental problem of threshold cryptography, a problem of secure sharing of a secret. A secret sharing scheme allows one to distribute a piece of secret information among several servers in a way that meets the following requirements: (1) no group of corrupt servers (smaller than a given threshold) can figure out what the secret is, even if they cooperate; (2) when it becomes necessary that the secret information be reconstructed, a large enough number of servers (a number larger than the above threshold) can always do it.

A very useful extension of secret sharing is function sharing. Its main idea is that a highly sensitive operation, such as decryption or signing, can be performed by a group of cooperating servers in such a way that no minority of servers is able to perform this operation by themselves, nor would they be able to prevent the other servers from performing the operation when it is required.

A good example of an application whose security could be greatly improved with a threshold solution is a network Certification Authority, a trusted entity that certifies that a given public key corresponds to a given user. If we trust one server to perform this operation, then it is possible that as a result of just one break-in, no certificate can any longer be trusted. Thus it is a good idea to distribute the functionality of the certification authority between many servers, so that an adversary would need to corrupt half of them before he can forge a certificate on some public key.

A certification authority is really a signature service: The public-key certificates it produces are signatures on messages that contain a description of some entity and its public key. Therefore, to implement such certification authority in a fault-tolerant threshold manner described above we need secure threshold signature schemes. There are many other function-sharing applications, including applications for distributed decryption. Threshold decryption schemes enable such operations as key recovery, organization's keys, fair sale of digital content in exchange for digital receipts; secure bidding, and secret election protocols.

1.3 Digital Signatures

Digital Signatures, not unlike the handwritten ones, are used to authenticate information: that is, to securely tie the contents of an electronic document to a signer (more precisely, to the signer's public key). Only the true signer should be able produce valid signatures, and anyone should be able to verify them in order to convince oneself that the signer indeed signed the document.

While many digital signature schemes have been proposed and a few are used in practice today, research into designing schemes that are more secure, more efficient, or have additional properties continues.

1.4 Property Testing

Probabilistic Property testing is a relatively new algorithmic technique which has been applied to testing NP-hard graph properties, testing algebraic and combinatorial functional properties such as linearity and monotonicity, and testing properties of classical linear algebra structures. These fast algorithmic methods have already yielded improvements in applied problems in the domain of program checking, web searching (where linear a;gebra plays a role), and proof checking. The field of property testing as it is known today was initiated by researchers Shafi Goldwasser, Oded Goldreich, and Dana Ron in the CIS group.

2 Progress Through December 1999

2.1 Quantum Computation

Grover demonstrated that quantum computers may outperform classical computers on general searching problems. Searching for an item on a list of N items requires O(N) steps classically, but Grover's quantum mechanical algorithm solves the problem in O(sqrt(N)) steps. Farhi and Gutmann subsequently 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 and Michael Sipser are in the process of completing a paper that gives a simpler Farhi-Gutmann type of analysis for the quantum counting algorithm

2.2 Threshold Cryptography

Our progress on ths problem during the last few months is summarized in a paper by Stanislaw Jarecki, and Anna Lysyanskaya titled "Adaptively secure threshold cryptosystems without erasures", submitted to Eurocrypt 1999.

This paper provides solutions for efficient threshold cryptosystems which are secure against adaptive adversaries even when the players cannot erase their local data. Specifically, it presents erasure-free adaptively-secure protocols for distributed key generation in discrete-log based schemes, for DSS and RSA signature generation, and for ElGamal decryption. Recently, [CGJKR99] introduced efficient adaptively-secure threshold cryptosystems whose security relies on the ability of the uncorrupted players to safely erase most of the secret data they produce during the protocol execution. However, secure erasure of data is hard to implement in practice: It requires specialized hardware and operating systems, and even then it would remain a costly operation.

This work builds directly on the protocols of [CGJKR99] for distributed generation of keys in discrete-log based cryptosystems, and for distributed generation of DSS and RSA signatures. By introducing a few subtle but crucial modifications to these protocols, and using novel techniques in the analysis of their security, thee need to recourse to erasures in these protocols is removed. These modifications actually reduce the complexity of the adaptively-secure RSA signature generation. However, in the key generation and DSS signature generation protocols, an increase in communication and computation complexity over the efficient [CGJKR99] protocols which is sublinear in the security parameter is incurred.

The sublinear increase in complexity is wholly due to the implementation of secure channels we use to make our protocols secure in the adaptive model. In contrast, the best currently known implementation of secure channels for general adaptively-secure erasure-free multiparty computation based on the DDH assumption, due to Beaver, creates an overhead which is linear in the security parameter.

[CGJKR99] Ran Canetti, Rosario Gennaro, Stanislaw Jarecki, Hugo Krawczyk, Tal Rabin. "Adaptive Security for Threshold Cryptosystems." Appeared in CRYPTO'99.

Stanislaw Jarecki, and Anna Lysyanskaya were partially supported by the NTT grant during Summer 99 and Fall 99.

2.3 Digital Signatures

Our progress for this period is written up in a paper by Silvio Micali, Kazuo Ohta and Leonid Reyzin, titled "Provable-Subgroup Signatures". Submission, 1999.

In this work, the notion of *Provable-Subgroup Signatures* (PS signatures), a generalization of multi-signature schemes is put forward. In essence, PS signatures enable any subgroup, *of a given group*, of potential signers, to sign efficiently a message so that the signature provably reveals the identities of the signers in to any verifier. (Thus, PS signatures keep the actual signers accountable for what they sign, which is desirable in several applications.) PS signatures remain secure even in the presence of a polynomial-time adversary who is allowed to corrupt arbitrary subsets of signers and mount (properly defined) chosen message attacks.

Our work provides two efficient implementations of PS signatures in the random-oracle model: one based on the Discrete Log problem, and the other based on a variant of the RSA. In both implementations, the resulting signatures are as compact as ordinary, single-signer signatures.

Kazuo Ohta is visiting the CIS group from NTT .

2.4 Property Testing

Work in this topic during the last six months is summarized in a paper by Yevgeniy Dodis, Oded Goldreich, Eric Lehman, Sofya Raskhodnikova, Dana Ron, and Alex Samorodnitsky, titled "Improved Testing Algorithms for Monotonicity", which appeared in the Proceedings of RANDOM, August 1999.

The paper presents improved algorithms for testing monotonicity of functions. Namely, given the ability to query an unknown function f, where Σ and Ξ are finite ordered sets, the test always accepts a monotone f, and rejects f with high probability if it is ϵ-far from being monotone (i.e., every monotone function differs from f on more than an ϵ fraction of the domain). For any $\epsilon > 0$, the query complexity of the test is $O((n/\epsilon) \cdot \log |\Sigma| \cdot \log |\Xi|)$.

The previous best known bound was $\tilde{O}((n^2/\epsilon) \cdot |\Sigma|^2 \cdot |\Xi|)$ by Shafi Goldwasser, Oded Goldreich, Eric Lehman, Dana Ron, and Alex Samorodnitsky.

We also present an alternative test for the boolean range Ξ whose query complexity $O(n^2/\epsilon^2)$ is independent of alphabet size $|\Sigma|$.

Eric Lehman was partially supported by the NTT grant during Fall99.

3 Plans for the Next Six Months

3.1 Quantum Computation

We are still attempting to look at nonrelativizing techniques that may apply to quantum computers. Such techniques offer the only hope of beating Grover's sqrt(N) bound, because Grover's result has been shown to be optimal if relativizability is allowed. We have some ideas for algorithms that may be able to improve upon Grover's result, but we have not been able to analyze them yet.

3.2 Cryptography

We are continuing to pursue research on all fronts, theoretical and cryptography. In particular, there are some on going projects on the All-Or-Nothing-transform (AONT) introduced by Rivest, and on question in concurrent protocol security.