abstract

COMPUTER SCIENCE/DISCRETE MATH II
Topic:

Speaker:

Affiliation:

Date:

Time/Room:

Current constructions of cryptographic primitives typically involve a large multiplicative computational overhead that grows with the desired level of security. We explore the possibility of implementing cryptographic primitives (such as encryption, authentication, signatures, or secure two-party computation) while incurring only a *constant* computational overhead compared to insecure implementations of the same tasks.

We obtain affirmative answers to this question for most central cryptographic primitives under plausible, albeit nonstandard, intractability assumptions.

* We start by showing that pairwise-independent hash functions can be computed by linear-size circuits, disproving a conjecture of Mansour, Nisan, and Tiwari (STOC 1990). This construction, which is used by subsequent cryptographic constructions, does not rely on any unproven assumptions and is of independent interest.

* Under an intractability assumption that generalizes a previous assumption of Alekhnovich (FOCS 2003), we get (public and private key) encryption schemes and MAC's that can be implemented by circuits whose size is a constant multiple of the length of the message to be encrypted or authenticated. A polynomial-time adversary attacking these schemes can only gain a negligible advantage in the length of the message. Under an exponentially strong version of the previous assumption (or exponentially strong versions of other related assumptions), we get signature schemes of a similar complexity.

* Assuming the existence of pseudorandom generators in NC-0 with polynomial stretch together with the existence of an (arbitrary) oblivious transfer protocol, we get similar results

for the seemingly very complex task of secure two-party computation. More concretely, we get general protocols for secure two-party computation in the semi-honest model in which the two parties can be implemented by circuits whose size is a constant multiple of the size s of the circuit to be evaluated. In the malicious model, we get protocols whose *communication complexity* is a constant multiple of s and whose computational complexity is slightly super-linear in s. For natural relaxations of security in the malicious model that are still meaningful in practice, we can also keep the computational complexity linear in s. These results extend to the case of a constant number of parties, where an arbitrary subset of the parties can be corrupted.

Our protocols rely on non-black-box techniques, and suggest the intriguing possibility that the ultimate efficiency in cryptographic protocols can be obtained via such techniques.

This is joint work with Yuval Ishai, Eyal Kushilevitz, and Rafail Ostrovsky, and will appear at STOC 2008.