Foundations of Cryptography

Take Home Exam
Date Due: March 11th 2007 – Strict Deadline

- **Instructions:** Questions 1 and 2 are compulsory. In addition, answer four more questions from 3–8 (six altogether).
- **Restrictions:** You should work on the exam alone. You can use your notes from class and the handouts. You should not use any other material while solving this exam.

1. Given a family $\mathcal{F}$ of pseudorandom permutations over $\{0,1\}^n$, we would like to construct (directly) a family of pseudorandom permutations over $\{0,1\}^{10n}$. For any $f_1, \ldots, f_{10} \in \mathcal{F}$ we define $h_{f_1,\ldots,f_{10}} : \{0,1\}^{10n} \to \{0,1\}^{10n}$ by $h_{f_1,\ldots,f_{10}}(x) = (f_1(x_1), \ldots, f_{10}(x_{10}))$, where $x = x_1 \cdots x_{10} \in \{0,1\}^{10n}$. Let $\mathcal{H} = \{h_{f_1,\ldots,f_{10}}\}_{f_1,\ldots,f_{10} \in \mathcal{F}}$.

   (a) Show that $\mathcal{H}$ is not a family of pseudorandom permutations.

   (b) Show that $\mathcal{H}$ is pseudorandom against a *random* attack.

2. A hybrid encryption scheme uses public-key encryption to encrypt a random symmetric key, and then proceeds to encrypt the message with that symmetric key. The receiver decrypts the symmetric key using the public-key encryption scheme and then uses the recovered symmetric key to decrypt the message.

   (a) Let $(\text{KeyGen}, \mathcal{E}, \mathcal{D})$ be a semantically secure public-key encryption scheme against a chosen-plaintext attack, and G a pseudorandom generator. Show that the following public-key encryption scheme in semantically secure against a chosen-plaintext attack:

   Key generation: The same as KeyGen.

   Encryption: Given a public key pk, choose a random seed s for the pseudorandom generator G and output $(\mathcal{E}_{pk}(s), G(s) \oplus m)$.

   Decryption: Given a secret key sk and a ciphertext (c_1, c_2), output $m = G(\mathcal{D}_{sk}(c_1)) \oplus c_2$.

   (b) Suppose that $(\text{KeyGen}, \mathcal{E}, \mathcal{D})$ in the above construction is CCA2-non-malleable. Is the resulting scheme CCA2-non-malleable as well?

   (c) Bonus: Suppose that $(\text{KeyGen}, \mathcal{E}, \mathcal{D})$ in the above construction is CCA1-non-malleable. What can you say about the security of the resulting scheme?

   (d) Bonus: Suppose that $(\text{KeyGen}, \mathcal{E}, \mathcal{D})$ in the above construction is CCA1-semantically-secure. What can you say about the security of the resulting scheme?

3. Consider the following commitment scheme:

   Public parameters: Alice and Bob share a prime number p, and two randomly chosen generators g and h of a subgroup of $\mathbb{Z}_p^*$ of prime order q.

   Commitment: In order to commit to $x \in \mathbb{Z}_q$, Alice picks a random $r \in \mathbb{Z}_q$ and sends $c = g^x h^r \bmod p$ to Bob.

   Reveal: Alice sends (x, r) to Bob, who verifies that $c = g^x h^r \bmod p$.

(a) Under the discrete-log intractability assumption, prove that the scheme is perfectly-hiding and computationally-binding.

(b) Suppose that Alice and Bob execute the following protocol:

- Alice commits to a value v_A by sending Bob a commitment C_A.
- Bob commits to a value v_B by sending Alice a commitment C_B.
- Alice opens C_A to reveal v_A to Bob.
- Bob opens C_B to reveal v_B to Alice.

Show that if the protocol is instantiated with the above commitment scheme, then Bob can always ensure that $v_B = 6v_A + 7$.

4. Assuming that one-way functions exist, prove that for all sufficiently large n there exist two polynomially-samplable distributions D_n and D'_n over n-vertex graphs, such that:

- D_n is a distribution over Hamiltonian graphs.
- D'_n is statistically-close to a distribution over non-hamiltonian graphs (I.e., there exists a distribution over non-hamiltonian graphs whose statistical distance to D'_n is negligible in n).
- D_n and D'_n are computationally indistinguishable.

5. Transform $\mathcal{H}$ from Question 1 into a family of pseudorandom permutations by using only "combinatorial objects".

6. **Strongly unforgeable signature schemes.** A signature scheme (KeyGen, Sign, Verify) is (t, q, ϵ)-strongly-unforgeable if any adversary that runs in time t and receives signatures on q adaptively chosen messages cannot produce a *new* signature on one of these q messages or on any other message, except with probability at most ϵ. A signature scheme is strongly-unforgeable if it is (t, q, ϵ)-strongly-unforgeable for all polynomials t and q, and for ϵ which is negligible in the security parameter.

(a) Suggest a strongly-unforgeable one-time signature scheme (i.e., $q = 1$) from any universal one-way hash function.
(Hint: modify the one-time signature scheme seen in class.)

(b) Let (KeyGen, Sign, Verify) be an existentially unforgeable (under an adaptive chosen message attack) signature scheme, and let $(KG_{OTS}, S_{OTS}, V_{OTS})$ be a strongly-unforgeable one-time signature scheme. Prove that the following scheme is strongly-unforgeable:

Key generation: The same as KeyGen.

Signature: On input private key sk and a message m, run KG_{OTS} to obtain a key pair (sk_{OTS}, vk_{OTS}) for the one-time signature, and output $\sigma = (\sigma_1, \sigma_2, vk_{OTS})$ where $\sigma_1 = \text{Sign}(sk, vk_{OTS})$ and $\sigma_2 = S_{OTS}(sk_{OTS}, m||\sigma_1)$.

Verification: On input public key pk, a message m and a signature $\sigma = (\sigma_1, \sigma_2, vk_{OTS})$, accept if and only if both $\text{Verify}(pk, \sigma_1, vk_{OTS})$ and $V_{OTS}(vk_{OTS}, \sigma_2, m||\sigma_1)$ accept.

(c) Bonus: Show that if the one-time signature scheme $(KG_{OTS}, S_{OTS}, V_{OTS})$ is not strongly unforgeable, then the resulting construction may not be strongly unforgeable.

7. Let D_n be an ensemble of distributions. A distribution D_n generates two strings (x, h) where $x \in \{0,1\}^n$ and $h \in \{0,1\}^m$ for some m polynomial in n. It is best to think of h as some "hint" about x. We are interested in compressing the string x (using hint h) in a manner that will allow reconstructing x without h. We say that the ensemble can be $k(n)$-compressed if there

exist two polynomial-time computable functions C and R such that for any (x, h) in the support of D_n it holds that $R(C(x, h)) = x$, and the expected length (over D_n) of $C(x)$ is at most $k(n)$. We say that D_n has non-trivial compression if it is $k(n)$-compressible for $k(n) \leq n - 1/p(n)$ for some polynomial $p(\cdot)$.

(a) Let $G : \{0,1\}^* \mapsto \{0,1\}^*$ be a pseudo-random generator. Show that the distribution $(G(s), \cdot)$ where s is a random seed does not have any non-trivial compression.

(b) Suppose that a verifier is given a string $y = G(s)$ and then engages in a zero-knowledge proof that y is in the domain of G. Can the verifier now come up with a non-trivial compression of y (expectation over random seed s)?

8. Our goal is to construct a 1-out-of-n OT protocol (secure against honest-but-curious adversaries) from any 1-out-2 OT protocol (secure in the same sense). Consider the following protocol:

- The sender has input $X_1, \ldots, X_N$, where $N = 2^n$, and the chooser has input $1 \leq I^* \leq N$.
- The sender prepares n pairs of random strings $(W_1^0, W_1^1), \ldots, (W_n^0, W_n^1)$, and for every $1 \leq I \leq N$ sets $Y_I = X_I \bigoplus_{j=1}^n W_j^{i_j}$ where $i_1 \cdots i_n$ is the binary representation of I. The strings $Y_1, \ldots, Y_N$ are sent to the chooser.
- For every $1 \leq j \leq n$, the parties execute a 1-out-of-2 OT protocol on the strings (W_j^0, W_j^1) in which the chooser wishes to learn $W_j^{i_j^*}$, where $i_1^* \cdots i_n^*$ is the binary representation of I^*.
- The chooser reconstructs $X_{I^*} = Y_{I^*} \bigoplus_{j=1}^n W_j^{i_j^*}$.

(a) Show that this is NOT a good protocol for 1-out-of-n OT (no matter what 1-out-of-2 OT protocol is used).

(b) Consider now a similar protocol, except that the masking of the X_I's is done differently. Let F_S be a pseudorandom function and treat the W_j^b's as keys to the function. Let $Y_I = X_I \bigoplus_{j=1}^n F_{W_j^{i_j}}(I)$. The rest of the protocol is as before, except that now the chooser reconstructs X_{I^*} by computing $Y_{I^*} \bigoplus_{j=1}^n F_{W_j^{i_j^*}}(I)$. Prove that this is a good 1-out-of-n protocol.

(c) Bonus: Devise a 1-out-of-n OT protocol with the same performance (i.e., $\log N$ executions of 1-out-of-2) based on the OT protocol seen in class.