CS-7880 Graduate Cryptography March 18, 2025

Problem Set 5

Lecturer: Daniel Wichs *Due:*

Problem 1 (Circular Security) 15pts

In this problem, we consider a security property (of encryption schemes) called **circular security**: namely, security even when an adversary is given an encryption of the secret key sk.

- First, consider public-key encryption. A public key encryption scheme $(\mathrm{Gen}, \mathrm{Enc}, \mathrm{Dec})$ is said to be *circularly secure* if any p.p.t. algorithm A wins the following game (interacting with a challenger) with probability at most $\frac{1}{2} + \mathrm{negl}(n)$:

  1. The challenger samples $(pk, sk) \leftarrow \mathrm{Gen}(1^n)$ and sends (pk, c^*) to A, where $c^* \leftarrow \mathrm{Enc}(pk, sk)$ is a ciphertext where the message is the secret key.
  2. A sends two messages (m_0, m_1) to the challenger.
  3. The challenger selects $b \leftarrow_R \{0,1\}$ and sends $c_b \leftarrow \mathrm{Enc}(pk, m_b)$ to A.
  4. A outputs a bit $\tilde{b}$. We say that A wins if $\tilde{b} = b$.

  It turns out that not every IND-CPA secure public key encryption (i.e., the standard notion of security we defined in class) is also circularly secure. Construct a public-key encryption scheme which is IND-CPA secure but not circularly secure, relying only on the existence of public-key encryption schemes. Prove that your scheme is IND-CPA secure but not circularly secure.

- Next consider symmetric-key encryption. A secret key encryption scheme $(\mathrm{Gen}, \mathrm{Enc}, \mathrm{Dec})$ is said to be *circularly secure* if for all p.p.t. oracle algorithms $A^{O(\cdot)}$, we have that

$$\left| \Pr_{sk \leftarrow_R \mathrm{Gen}(1^n)} \left[A^{\mathrm{Left}_{\mathrm{sk}}(\cdot)}(c^*) = 1 \right] - \Pr_{sk \leftarrow_R \mathrm{Gen}(1^n)} \left[A^{\mathrm{Right}_{\mathrm{sk}}(\cdot)}(c^*) = 1 \right] \right| \leq \mathrm{negl}(n), \qquad (1)$$

  where $c^* \leftarrow \mathrm{Enc}(sk, sk)$, $\mathrm{Left}_{\mathrm{sk}}(\cdot)$ is an oracle that on input (m_L, m_R) outputs an encryption of m_L, and $\mathrm{Right}_{\mathrm{sk}}(\cdot)$ is an oracle that on input (m_L, m_R) outputs an encryption of m_R.

  show that a variant of the LWE-based secret key encryption we saw in class does satisfy circular security (under an LWE-like assumption). In particular, we consider a variant of the LWE problem where the secret s is a uniformly random binary string.

  Definition 3 (LWE* Assumption): The LWE* assumption with error distribution χ states that the following two distributions are computationally indistinguishable:

$$\{s \leftarrow_R \{0,1\}^n, A \leftarrow_R \mathbb{Z}_q^{n \times m}, e \leftarrow \chi^m : (A, s^T A + e^T)\} \approx_c \{A \leftarrow_R \mathbb{Z}_q^{n \times m}, b \leftarrow_R \mathbb{Z}_q^m : (A, b^T)\}.$$

Under the LWE* assumption (with m any polynomial in n, and with error distribution χ), prove that the n-bit encryption scheme defined by

$$\text{Enc}(s \in \{0,1\}^n, m \in \{0,1\}^m; R \leftarrow_R \mathbb{Z}_q^{n \times m}, e \leftarrow \chi^m) = (R, s^T R + e^T + \lfloor \frac{q}{2} \rfloor m^T)$$

is a circularly secure secret key encryption scheme.

Hint: Show that it is possible to generate an encryption of s given only an encryption of 0.

Problem 2 (Collision Resistance from LWE) 10 pts

Consider the hash function $H_A(r) = rA$ where $A \leftarrow \mathbb{Z}_q^{m \times n}$ is the hash key, $r \in \{0,1\}^m$ is the input (interpreted as a row vector), and all operations are performed over $\mathbb{Z}_q$. This function is compressing when $m > n \log q$. Show that assuming the hardness of the decision LWE assumption with error chosen from the interval $[-B, B]$, where $mB < q/4$, the given hash function is collision resistant. Your proof should proceed in 2 steps:

- Show that if H_A is not collision resistant then there exists a p.p.t. $\mathcal{A}$ such that

$$\Pr[r^* A = 0, r^* \in \{-1, 0, 1\}^m, r^* \neq 0 \ : \ A \leftarrow \mathbb{Z}_q^{m \times n}, r^* \leftarrow \mathcal{A}(A)]$$

  is non-negligible.

- Show that the above $\mathcal{A}$ can be used to break decision LWE.

Problem 3 (Signatures from OWFs) 10 pts

A weak collision-resistant hash function family (also called universal one-way hash in the literature) $H_s \ : \ \{0,1\}^* \to \{0,1\}^n$ with seed $s \in \{0,1\}^n$ ensures that every p.p.t. $\mathcal{A}$ has at most a negligible chance of winning the following game:

- $\mathcal{A}$ chooses $x \in \{0,1\}^*$.
- The challenger chooses a random seed $s \leftarrow \{0,1\}^n$.
- $\mathcal{A}$ gets s and chooses $x' \in \{0,1\}^*$.
  It wins if $x \neq x'$ and $H_s(x) = H_s(x')$.

This is weaker than collision resistance since one of the values x involved in the collision has to be chosen by the adversary before it knows the seed s. It turns out that such weak collision-resistant hashing can be constructed from just one-way function (in contrast, we do not believe we can construct standard collision resistant hashing from one-way functions).

Show how to adapt the signature scheme based on collision-resistance that we saw in class to only rely on weak collision resistance. Make sure your signature scheme is capable of signing arbitrarily long messages. You do not have to write the scheme in full, but discuss the differences from the one in class, both in terms of the scheme construction and the security analysis.

Problem 4 (Commitments) 15 pts

We used commitments to construct ZK proofs. A commitment scheme allows Alice to "commit" herself to some message m by giving Bob some value $c = \mathsf{Commit}(m, r)$ generated using randomness r. Bob should not learn anything about m given the commitment c. Later she can "open" the commitment by giving (m, r) to Bob to convince him that m was the value she committed herself to. We will assume that $m \in \{0,1\}$ is just a single bit – we can always extend this to a longer message by committing one bit at a time.

Formally, a commitment is a function $\mathsf{Commit} : \{0,1\} \times \{0,1\}^* \to \{0,1\}^*$ that should satisfy the following properties.

- Hiding: The commitments to 0 and 1 are computationally indistinguishable $\mathsf{Commit}(0, U_n) \approx \mathsf{Commit}(1, U_n)$ where U_n denotes the uniform distribution over $\{0,1\}^n$.
- Binding: For all PPT adversaries A, we have

$$\Pr[\mathsf{Commit}(0, r) = \mathsf{Commit}(1, r') : (r, r') \leftarrow A(1^n)] = \mathsf{negl}(n).$$

A. Let f be a one-way permutation with a hardcore predicate h. Show that $\mathsf{Commit}(b, r) = (f(r), h(r) \oplus b)$ is a secure commitment scheme where the biding property holds even if the adversary is computationally unbounded.

B. You will now show how to construct a generalized notion of such commitments, which we'll call seeded commitments, from one-way functions. A seeded commitment also contains a seed s and we define the commitment function as $\mathsf{Commit}_s(m, r)$ which takes the seed s as an input. We think of Bob as generating the seed s and therefore we want hiding to hold even if s is chosen maliciously. On the other hand, we want binding to hold when s is chosen randomly. Give a formal definition of the hiding and binding properties for seeded commitments to capture this intuitive description.

Consider the following scheme: Let G be a PRG with $2n$-bit stretch, so that for $|r| = n$, $|G(r)| = 3n$. Let $s \leftarrow \{0,1\}^{3n}$ be a random string of length $3n$. Define $\mathsf{Commit}_s(m, r)$ to be $G(r)$ if $m = 0$ and $G(r) \oplus s$ if $m = 1$. Show that this scheme satisfies the definition of seeded commitments, and that the binding property holds even if the adversary is computationally unbounded.

(hint: to argue binding, need to show that with high probability over a random s, there does not exist any pair r, r' such that $G(r) = G(r') \oplus s$. Use the union bound.)

C. In the schemes from part A and B, the hiding property holds when the adversary is computationally bounded but binding holds even if the adversary is computationally unbounded. We could ask whether the reverse is also possible.

Consider the following commitment scheme based on the discrete-logarithm assumption. The seed s consists of $s = (\mathbb{G}, q, g, h)$ where $(\mathbb{G}, q, g) \leftarrow \mathsf{GroupGen}(1^n)$ is a description of a cyclic group $\mathbb{G}$ of prime order q with generator g, and $h \leftarrow \mathbb{G}$ is a random group element. We define $\mathsf{Commit}_s(m; r) = g^m h^r$ where $m \in \{0,1\}$, $r \in \mathbb{Z}_q$. (We're changing the syntax a little so that the randomness is uniform over $\mathbb{Z}_q$ rather than $\{0,1\}^n$. We could also allow m to come from all of $\mathbb{Z}_q$ rather than just $\{0,1\}$ but let's stick with 1-bit messages to keep the syntax and the definitions consistent).

Show that the above scheme satisfies perfect hiding: for any s the distributions $\mathsf{Commit}_s(0, r) \equiv \mathsf{Commit}_s(1, r)$ are identical over the choice of a random r. Show that, under the discrete logarithm assumption, the scheme is binding when s is chosen randomly as specified above.