Problem Set 1

A glossary of formal definitions is provided for reference on the following page. Hints are available at the end of the document, but you are encouraged to try each problem first.

The first two problems will (re)introduce you to the notion of a *hybrid argument*, which is ubiquitous in cryptography. If you are unfamiliar, there are many resources on online, e.g. https://www.cs.columbia.edu/ tal/4261/F14/hybrid.pdf.

1. Assume the existence of a statistically hiding, collapse-binding commitment scheme $\mathsf{Com}_{\mathsf{ck}}(m, \omega)$ for one bit messages. Construct a statistically hiding, collapse-binding commitment scheme for arbitrary length messages ($\mathsf{poly}(\lambda)$ bits).

2. An efficiently computable function[1] h is *collapsing* if an adversary that outputs a superposition of valid input-output pairs $(x, h(x))$ can't distinguish whether:

   - the challenger measures x, or
   - the challenger measures $h(x)$.

   (See the glossary for a formal definition)

   Assume the existence of a collapsing hash function $h : \{0,1\}^{2\lambda} \to \{0,1\}^{\lambda}$. Construct a collapsing hash function $h' : \{0,1\}^{4\lambda} \to \{0,1\}^{\lambda}$.

3. Suppose $\mathsf{Com}_{\mathsf{ck}}(m, \omega)$ is a collapse-binding commitment for one-bit messages m. Prove that Com also satisfies the following definitions:

   (a) Classical-style binding: no QPT adversary given a random commitment key ck can output (ω_0, ω_1) satisfying $\mathsf{Com}_{\mathsf{ck}}(0, \omega_0) = \mathsf{Com}_{\mathsf{ck}}(1, \omega_1)$.

   (b) Sum-binding: after sending the commitment c, the probability that any QPT adversary can provide a valid ω_b given a random $b \leftarrow \{0,1\}$ is at most $1/2 + \mathsf{negl}(\lambda)$. This notion was implicitly defined in class, and the solution to this problem should follow from reviewing the lecture notes.

   Now, suppose $\mathsf{Com}_{\mathsf{ck}}(m, \omega)$ is a collapse-binding commitment for $\mathsf{poly}(\lambda)$-bit messages m. Prove that Com also satisfies the following definition:

   (c) Collapse-binding w.r.t. any efficient function f: this is identical to collapse-binding, except that when $b = 1$, the challenger computes an efficient function f on the message and only measures $f(m)$. (Note that collapse-binding corresponds to $f(m) = m$.)

4. (Challenging) Suppose $h : \{0,1\}^{\lambda+1} \to \{0,1\}^{\lambda}$ has the following properties:

   (a) for all $x \in \{0,1\}^{\lambda}$, there is exactly one $x' \in \{0,1\}^{\lambda}$ satisfying $h(0, x) = h(1, x')$, and

   (b) no QPT adversary can output x, x' such that $h(0, x) = h(1, x')$.[2]

[1] Technically, we mean hash function *family*; see the definition glossary for this formalism.

[2] A hash function with this pair of properties is sometimes referred to as "claw-free".

Prove that h is a collapsing hash function.

5. (Challenging) Suppose $\mathsf{Com}_{\mathsf{ck}}(m, \omega)$ is a commitment scheme for one-bit messages satisfying sum-binding. Prove that $\mathsf{Com}_{\mathsf{ck}}$ must be collapse-binding. Combined with $(2b)$, this establishes that sum binding and collapse binding are equivalent notions for bit commitments.

6. Suppose you are given a magical oracle Cl that can clone any quantum state, i.e., for any $|\psi\rangle$, $\mathsf{Cl}(|\psi\rangle) \to |\psi\rangle \otimes |\psi\rangle$. Prove that there do not exist collision-resistant hash functions in this world.

7. (Optional bonus problem)

   (a) The definition of sum binding applies to *bit* commitments. Consider the following definition that attempts to extend sum-binding to commitments for n-bit messages: after sending the commitment c, the probability that any QPT adversary given a random $m \leftarrow \{0,1\}^n$, can output an opening ω_m satisfying $c = \mathsf{Com}_{\mathsf{ck}}(m, \omega_m)$ is at most $1/2^n + \mathsf{negl}(\lambda)$. It turns out that this definition captures binding when $n = O(\log \lambda)$ but fails for larger n. Explain why.

   (b) (Challenging) Propose a sum-binding-style definition that works for any $n = \mathsf{poly}(\lambda)$ length messages and prove that it is equivalent to collapse-binding. Your security game should only involve classical communication.

8. (Optional bonus problem) In class, we proved that lossy functions imply collapsing hash functions. The goal of this problem is to prove that, in turn, collapsing hash functions imply (statistically-hiding) collapse-binding bit commitments. For a collapsing hash function $H : \{0,1\}^n \to \{0,1\}^{n/2}$, consider the bit commitment scheme that works as follows: to commit to a bit m, sample random $x, r \leftarrow \{0,1\}^n$ and output $(H(x), r, m \oplus \langle x, r\rangle)$. Prove that this scheme is collapse-binding.

9. (Optional bonus problem) In this problem, you will construct lossy functions assuming that the decisional Learning with Errors (LWE) problem is hard. Concretely, the decisional LWE assumption, denoted $\mathsf{LWE}_{n,m,q,\chi}$, is parameterized by $\mathsf{poly}(\lambda)$-size parameters $n(\lambda), m(\lambda)$ (where $n \leq m$), a modulus $q(\lambda)$ (which may be superpolynomial), and a probability distribution $\chi(\lambda)$ with bounded support $\beta \ll q$. It states that for random $A \leftarrow \mathbb{Z}_q^{m \times n}, s \leftarrow \mathbb{Z}_q^n, e \leftarrow \chi^m$, and $u \leftarrow \mathbb{Z}_q^m$, then

$$(A, As + e) \approx_c (A, u),$$

   where $\approx_c$ denotes computational indistinguishability. For additional background on the LWE problem, see, e.g., https://people.csail.mit.edu/vinodv/CS294/lecture1.pdf.

   (a) Let $n'(\lambda) < n(\lambda)$. Show that the $\mathsf{LWE}_{n',m,q,\chi}$ assumption implies that

$$BC + E \approx_c A.$$

   where $B \leftarrow \mathbb{Z}_q^{m \times n'}, C \leftarrow \mathbb{Z}_q^{n' \times n}, E \leftarrow \chi^{m \times n}$, and $A \leftarrow \mathbb{Z}_q^{m \times n}$. Notice that the left-hand-side corresponds to a rank $n' < n$ matrix plus a "noise" term E, while the right-hand-side corresponds to a random matrix.

(b) Define the following "rounding" operation

$$\lfloor\cdot\rfloor_p : \mathbb{Z}_q \to \mathbb{Z}_p : x \to \lfloor(p/q)x\rfloor.$$

In words, the function $\lfloor\cdot\rfloor_p$ can be computed by splitting the range $\{0,1,\ldots,q-1\}$ into p intervals (each of size each of size $\leq q/p$) $\{0,1,\ldots,\lfloor q/p\rfloor-1\},\{\lfloor q/p\rfloor,\ldots,\lfloor 2q/p\rfloor-1\},\ldots,$ and mapping any input in the ith interval (indexing from 0) to $i \in \mathbb{Z}_p$.

Suppose $n(\lambda)$ and $m(\lambda)$ satisfy $m > n^2$. Suppose $p(\lambda)$ and $q(\lambda)$ are such that q/p is superpolynomial and $m \log p > \lambda n^2$. Let $\chi = \chi(\lambda)$ be a β-bounded distribution for some polynomial $\beta(\lambda)$.

Define a keyed hash function $H_A : \{0,1\}^n \to \mathbb{Z}_p^m$ as

$$H_A(x) := \lfloor Ax\rfloor_p,$$

where A is the hash key (i.e., the function description). Prove that H_A is a lossy function (Definition 0.7). In particular, prove both of the following:

i. When A is sampled as $A \leftarrow \mathbb{Z}_q^{m\times n}$, H_A is injective with all but negligible probability.

ii. When A is sampled as $BC + E$ as described above, H_A is $(n - n'\log q)$-lossy with all but negligible probability.

Glossary of Cryptographic Definitions

Cryptographic primitives are typically parameterized by a security parameter λ that roughly corresponds to how "hard" it is to break security. An efficient quantum polynomial-time (QPT) adversary can be described by a family of $\mathsf{poly}(\lambda)$-size quantum circuits $\{\mathcal{A}_\lambda\}_\lambda$.

Remark 0.1 (quantum advice). *Modeling an adversary as a family of quantum circuits $\{\mathcal{A}_\lambda\}_\lambda$ captures attacks with $\mathsf{poly}(\lambda)$-size classical advice (e.g., the advice can be part of the circuit description). While not necessary for this problem set, it is also possible to consider adversaries with explicit quantum advice, i.e., $\{\mathcal{A}_\lambda, \rho_\lambda\}_\lambda$ where ρ_λ is an arbitrary $\mathsf{poly}(\lambda)$-size quantum state.*

A function $f : \mathbb{N} \to \mathbb{R}$ is *negligible* if it decays faster than the inverse of any polynomial, i.e., for any $c > 0$, there exists an integer λ_c such that for all $\lambda > \lambda_c$, $|f(\lambda)| < 1/\lambda^c$. In the following definitions, security requires that any $\mathsf{poly}(\lambda)$-time quantum adversary wins the associated security game with probability at most $\mathsf{negl}(\lambda)$ (for search games), or $1/2 + \mathsf{negl}(\lambda)$ (for decision games).

Cryptographic hash functions

Definition 0.2 (Collision resistant hash functions). *A family of hash functions $\{H_\lambda\}_\lambda$ is (post-quantum)* collision-resistant *if for any quantum polynomial-time (henceforth, QPT) adversary $\{\mathcal{A}_\lambda\}_\lambda$,*

$$\Pr\left[\begin{array}{c} h(x_0) = h(x_1) \text{ and } x_0 \neq x_1 \end{array} : \begin{array}{r} h \leftarrow H_\lambda \\ (x_0, x_1) \leftarrow \mathcal{A}_\lambda(h) \end{array}\right] = \mathsf{negl}(\lambda).$$

Definition 0.3 (Collapsing hash functions). *A family of hash functions $\{H_\lambda\}_\lambda$ is* collapsing *if for any QPT adversary $\{\mathcal{A}_\lambda\}_\lambda$, the following experiment outputs 1 with probability at most $1/2 + \mathsf{negl}(\lambda)$.*

1. *The challenger samples $h \leftarrow H_\lambda$ and sends h to $\mathcal{A}_\lambda$.*
2. *$\mathcal{A}_\lambda$ outputs a state on register X, which is supported on the domain of h.*
3. *The challenger computes the function $h(\cdot)$ in superposition, and measures the output, obtaining a string y. Next, the challenger samples $b \leftarrow \{0, 1\}$.*
   - *If $b = 0$, it returns the state on register X to the adversary.*
   - *If $b = 1$, it measures X in the standard basis and then returns the resulting state to the adversary.*
4. *$\mathcal{A}_\lambda$ outputs a bit b'. The experiment outputs 1 if $b = b'$.*

Commitment schemes

A non-interactive commitment scheme for a message space M consists of a pair of algorithms $(\mathsf{Gen}, \mathsf{Com})$.

- $\mathsf{Gen}(1^\lambda)$ is a randomized algorithm that outputs a commitment key ck.
- $\mathsf{Com}_{\mathsf{ck}}(1^\lambda, m, \omega)$ is a deterministic algorithm that takes as input a message $m \in M$ and a string $\omega \in \{0,1\}^\lambda$ corresponding to the opening/commitment randomness. (We typically drop the 1^λ when it is clear from context.)

To commit to a message $m \in M$, the sender samples a string random $\omega \leftarrow \{0,1\}^\lambda$ and outputs $c = \mathsf{Com}_{\mathsf{ck}}(1^\lambda, m, \omega)$. To open, the sender reveals (m, ω), and the receiver verifies by checking that $\mathsf{Com}_{\mathsf{ck}}(1^\lambda, m, \omega) = c$.

Definition 0.4 (Classical-style binding). *A non-interactive commitment scheme* $(\mathsf{Gen}, \mathsf{Com})$ *satisfies (post-quantum) classical-style binding if for any QPT adversary* $\{\mathcal{A}_\lambda\}_\lambda$*,*

$$\Pr\left[\mathsf{Com}_{\mathsf{ck}}(m_0,\omega_0) = \mathsf{Com}_{\mathsf{ck}}(m_1,\omega_1) \text{ and } m_0 \neq m_1 : \begin{array}{r} \mathsf{ck} \leftarrow \mathsf{Gen}(1^\lambda) \\ (m_0,\omega_0,m_1,\omega_1) \leftarrow \mathcal{A}_\lambda(\mathsf{ck}) \end{array}\right] = \mathsf{negl}(\lambda).$$

Definition 0.5 (Collapse binding). *A non-interactive commitment scheme* $(\mathsf{Gen}, \mathsf{Com})$ *satisfies collapse binding if for any QPT adversary* $\{\mathcal{A}_\lambda\}_\lambda$*, the following experiment outputs 1 with probability at most* $1/2 + \mathsf{negl}(\lambda)$*.*

1. *The challenger samples* $\mathsf{ck} \leftarrow \mathsf{Gen}(1^\lambda)$ *and sends* ck *to* $\mathcal{A}_\lambda$*.*
2. $\mathcal{A}_\lambda$ *outputs a state on registers* (M, W)*, where* M *is supported on* M *and* W *is supported on* $\{0,1\}^\lambda$ *(i.e.,* W *holds* λ *qubits).*
3. *The challenger computes the function* $\mathsf{Com}_{\mathsf{ck}}(\cdot,\cdot)$ *in superposition on* (M, W) *and measures the output, obtaining a string* c*. Next, the challenger samples* $b \leftarrow \{0,1\}$*.*
   - *If* $b = 0$*, it returns the state on registers* (M, W) *to the adversary.*
   - *If* $b = 1$*, it measures* M *in the standard basis and then returns the resulting state on* (M, W) *to the adversary.*
4. $\mathcal{A}_\lambda$ *outputs a bit* b'*. The experiment outputs 1 if* $b = b'$*.*

The following definition only applies to bit commitment schemes (i.e., the message space M is $\{0,1\}$).

Definition 0.6 (Sum binding). *A non-interactive bit commitment scheme* $(\mathsf{Gen}, \mathsf{Com})$ *satisfies sum binding if for any QPT adversary* $\{\mathcal{A}_\lambda\}_\lambda$*, the following experiment outputs 1 with probability at most* $1/2 + \mathsf{negl}(\lambda)$*.*

- *The challenger samples* $\mathsf{ck} \leftarrow \mathsf{Gen}(1^\lambda)$ *and sends* ck *to* $\mathcal{A}_\lambda$*.*
- $\mathcal{A}_\lambda$ *outputs a commitment* c*.*
- *The challenger flips a coin* $b \leftarrow \{0,1\}$ *and sends* b *to* $\mathcal{A}_\lambda$*.*
- $\mathcal{A}_\lambda$ *outputs* ω*. The experiment outputs 1 if* $c = \mathsf{Com}_{\mathsf{ck}}(b, \omega)$*.*

Definition 0.7 (Lossy function). *An ℓ-lossy function is a keyed family of hash functions with two ways to sample the key: $k \leftarrow \mathsf{InjSamp}(1^\lambda)$ and $k \leftarrow \mathsf{LossySamp}(1^\lambda)$. Each key k defines a hash function $H_k : \{0,1\}^n \to \{0,1\}^m$. We require the following properties.*

- *With overwhelming probability over $k \leftarrow \mathsf{InjSamp}(1^\lambda)$, H_k is injective, meaning that $|\{y : \exists x \text{ s.t. } H_k(x) = y\}| = 2^n$.*
- *With overwhelming probability over $k \leftarrow \mathsf{LossySamp}(1^\lambda)$, H_k is ℓ-lossy, meaning that $|\{y : \exists x \text{ s.t. } H_k(x) = y\}| \leq 2^{n-\ell}$.*
- *For any QPT adversary $\{\mathcal{A}_\lambda\}_\lambda$,*

$$\left| \Pr_{k \leftarrow \mathsf{InjSamp}(1^\lambda)}[\mathcal{A}_\lambda(k) \to 1] - \Pr_{k \leftarrow \mathsf{LossySamp}(1^\lambda)}[\mathcal{A}_\lambda(k) \to 1] \right| = \mathsf{negl}(\lambda).$$

Hints

2. Define $h'(x_1, x_2) := h(h(x_1), h(x_2))$.

3. For part (a), consider what happens if a collapsing adversary could efficiently prepare a state of the form $\frac{1}{\sqrt{2}}(|0, \omega_0\rangle + |1, \omega_1\rangle)$. For part (c), consider a reduction that uses the partial-collapsing adversary $\mathcal{A}$ and performs the measurement f coherently on the message register output by $\mathcal{A}$.

4. You may find it helpful to prove the following information-theoretic lemma:

   Lemma 0.8. *Let D be a projector and let Π_0, Π_1 be orthogonal projectors (that is, $\Pi_0\Pi_1 = 0$). $|\psi\rangle$ be a state that $|\psi\rangle \in \mathsf{span}(\Pi_0 + \Pi_1)$. Prove that*

   $$\left\|\Pi_1 D \Pi_0 |\psi\rangle\right\|^2 + \left\|\Pi_0 D \Pi_1 |\psi\rangle\right\|^2 \geq \frac{1}{2}\left(\|D |\psi\rangle\|^2 - \left(\|D\Pi_0 |\psi\rangle\|^2 + \|D\Pi_1 |\psi\rangle\|^2\right)\right)^2.$$

   Notice that the RHS corresponds to the advantage that D has in distinguishing $|\psi\rangle$ from the mixed state $\sum_{b\in\{0,1\}} \Pi_b |\psi\rangle\langle\psi| \Pi_b$.

   A hint for proving the lemma: Begin by writing

   $$\|D|\psi\rangle\|^2 = \langle\psi| D |\psi\rangle = \langle\psi| (\Pi_0 + \Pi_1) D (\Pi_0 + \Pi_1) |\psi\rangle,$$

   where the second equality follows because $|\psi\rangle \in \mathsf{span}(\Pi_0 + \Pi_1)$.

5. This should be a very similar proof to your proof for problem 4.

8. For part (a), use a hybrid argument over the columns of C, E. For part (b.i), show via a union bound that the function $x \to \lceil (q/p) H_A(x) \rfloor$ is injective (with overwhelming probability over the sampling of A). For part (b.ii), show via a union bound that the size of $\{y : y = \lfloor BCx + Ex \rfloor_p\}_{x\in\{0,1\}^n}$ is no larger than the size of $\{z : z = BCx\}_{x\in\{0,1\}^n}$ (with overwhelming probability over the sampling of B).