MIT 6.5620/6.875/18.425: Foundations of Cryptography **Due: Nov. 18, 2026**

Problem Set 4: Solutions

Your name **Your Collaborators**

Total Number of Points: 52.

Instructions: Solutions should be typeset in LaTeX and submitted as PDF by 11:59:59 PM ET on the due date. You can collaborate on the psets in small groups (of up to three), but must write your solutions separately, in your own words. This means not copying from a board nor from your collaborator's computer.

We strongly recommend ***not*** using AI tools to solve your problem sets. You are responsible for understanding everything you submit and you must be able to explain and justify every part of your solution. In particular, the oral problem set review (10% of the course grade) will evaluate your understanding of the solutions you submit. In all cases, you must list your collaborators, human or machine.

Problem 1. Hash It Out. (14 points)

In this problem, we will explore four different definitions of hash functions. We say that a family of functions $\mathcal{H} = \{H_k : \{0,1\}^{2n} \to \{0,1\}^n\}_{k \in \{0,1\}^m}$ is a X hash function, for $\mathsf{X} \in \{$"universal", "weak collision resistant", "universal one-way"[1], "collision resistant"$\}$, if it has the following properties:

- **Efficiency**: Given $k \in \{0,1\}^m$ (where $m = \mathsf{poly}(n)$) and $x \in \{0,1\}^{2n}$, $H_k(x) \in \{0,1\}^n$ can be computed in $\mathsf{poly}(n)$ time.
- **Security**:
  - If $\mathsf{X} =$ "universal": There exists a negligible function negl such that for all $x_1 \neq x_2 \in \{0,1\}^{2n}$,
$$\mathbf{Pr}\left[k \leftarrow \{0,1\}^m : H_k(x_1) = H_k(x_2) \wedge x_1 \neq x_2\right] \leq \mathsf{negl}(n).$$
  - If $\mathsf{X} =$ "weak collision resistant": For all PPT adversaries $\mathcal{A}$, there exists a negligible function negl such that the following holds:
$$\mathbf{Pr}\left[(k, x_1) \leftarrow \{0,1\}^m \times \{0,1\}^{2n}; \mathcal{A}(k, x_1) \to x_2 : H_k(x_1) = H_k(x_2) \wedge x_1 \neq x_2\right] \leq \mathsf{negl}(n).$$
  - If $\mathsf{X} =$ "universal one-way": For all PPT adversaries $\mathcal{A}$, there exists a negligible function negl such that for all $x_1 \in \{0,1\}^{2n}$, the following holds:
$$\mathbf{Pr}\left[k \leftarrow \{0,1\}^m; \mathcal{A}(k, x_1) \to x_2 : H_k(x_1) = H_k(x_2) \wedge x_1 \neq x_2\right] \leq \mathsf{negl}(n).$$
  - If $\mathsf{X} =$ "collision resistant": For all PPT adversaries $\mathcal{A}$, there exists a negligible function negl such that the following holds:
$$\mathbf{Pr}\left[k \leftarrow \{0,1\}^m; \mathcal{A}(k) \to (x_1, x_2) : H_k(x_1) = H_k(x_2) \wedge x_1 \neq x_2\right] \leq \mathsf{negl}(n).$$

(a) (2 points) Construct a family of universal hash functions without using any cryptographic assumptions. *Hint: Consider multiplication by a random matrix* $A \leftarrow \{0,1\}^{n \times 2n}$.

Solution

Consider a key given by a random matrix $A \leftarrow \{0,1\}^{n \times 2n}$ (so $m = 2n^2$), and let $H_k(x) = Ax \pmod 2$. Efficiency is clear since this is a matrix-vector product. To see security, fix any $x_1 \neq x_2$. We have
$$\mathop{\mathbf{Pr}}_{A \leftarrow \{0,1\}^{n \times 2n}}[Ax_1 = Ax_2] = \mathop{\mathbf{Pr}}_{A \leftarrow \{0,1\}^{n \times 2n}}[Ax = 0]$$

[1] also called "targeted collision resistant"

for $x = x_1 - x_2 \neq 0$. Since $x \neq 0$, we know it has some non-zero coordinate $i^* \in [2n]$. Therefore, for each $j \in [n]$,

$$(Ax)_j = \sum_{i=1}^{2n} A_{ji}x_i = A_{ji^*} + \sum_{i \neq i^*} A_{ji}x_i.$$

Since A_{ji^*} is i.i.d. uniformly random, the probability that $(Ax)_j = 0$ for all $j \in [n]$ is 2^{-n}. This is negligible, as desired.

(b) (3 points) Show how to build a universal one-way hash function family from any weak collision resistant hash function family.

Solution

Let $\mathcal{H} = \{H_k : \{0,1\}^{2n} \to \{0,1\}^n\}_{k \in \{0,1\}^m}$ be a weak collision resistant hash function family. Define $\mathcal{H}' = \{H_{k'} : \{0,1\}^{2n} \to \{0,1\}^n\}_{k' \in \{0,1\}^{m+2n}}$ by decomposing $k' = (k, r) \in \{0,1\}^m \times \{0,1\}^{2n}$ and setting $H'_{k'}(x) = H_k(x \oplus r)$. This clearly satisfies efficiency, so it suffices to show UOWHF security.

Suppose for the sake of contradiction that there exists an $x'_1 \in \{0,1\}^{2n}$ and a PPT adversary A' such that

$$\Pr_{k' \leftarrow \{0,1\}^{m+2n}}[A'(k', x'_1) \to x'_2 : H'_{k'}(x'_1) = H'_{k'}(x'_2) \wedge x'_1 \neq x'_2] \geq 1/p(n)$$

infinitely often for some polynomial $p(n)$. We can define an adversary A for H as follows: On input (k, x_1) to A, let $r := x_1 \oplus x'_1$ and $k' := (k, r)$, and run $A'(k', x'_1)$ to get x'_2. Output $x_2 := x'_2 \oplus r$.

Clearly A is PPT if A' is. Moreover, in the weak collision resistance game, since x_1 is uniformly randomly generated independently of k, we know that the joint distribution $k' = (k, r)$ is uniformly random (by a one-time pad argument), so we can invoke the guarantee of A' that with $1/p(n)$ probability infinitely often,

$$H'_{k'}(x'_1) = H'_{k'}(x'_2) \wedge x'_1 \neq x'_2,$$

which, by construction of H', implies

$$H_k(x'_1 \oplus r) = H_k(x'_2 \oplus r) \wedge x'_1 \neq x'_2.$$

Since $x_1 = x'_1 \oplus r$ and $x_2 = x'_2 \oplus r$, it follows that $x_1 \neq x_2$ and $H_k(x_1) = H_k(x_2)$. Therefore, $\mathcal{H}$ is not a weak collision resistant hash function family, as desired.

(c) (4 points) Show that there exists a hash family that is universal one-way but not collision resistant. (You may assume the existence of universal one-way hash functions.)

Solution

Let $\mathcal{H} = \{H_k : \{0,1\}^{2n} \to \{0,1\}^n\}_{k \in \{0,1\}^m}$ be a UOWHF. Define $\mathcal{H}' = \{H'_{k'} : \{0,1\}^{2n} \to \{0,1\}^n\}_{k' \in \{0,1\}^{m+4n}}$ by parsing $k' = (k, x_1^*, x_2^*) \in \{0,1\}^m \times \{0,1\}^{2n} \times \{0,1\}^{2n}$ and defining

$$H'_{k'}(x') := \begin{cases} H_k(x_2^*) & \text{if } x' = x_1^*, \\ H_k(x') & \text{otherwise.} \end{cases}$$

Clearly, $\mathcal{H}'$ is efficiently computable. First, we show that $\mathcal{H}'$ is universal one-way. Suppose not for contradiction. Then, there is some $x'_1 \in \{0,1\}^{2n}$ and some PPT adversary A' such that

$$\Pr_{k' \leftarrow \{0,1\}^{m+2n}}[A'(k', x'_1) \to x'_2 : H'_{k'}(x'_1) = H'_{k'}(x'_2) \wedge x'_1 \neq x'_2] \geq 1/p(n),$$

for some polynomial $p(n)$ infinitely often. Equivalently,

$$\Pr_{k\leftarrow\{0,1\}^{m+2n},x_1^*,x_2^*\leftarrow\{0,1\}^{2n}}[A'(k,x_1^*,x_2^*,x_1')\to x_2' : H'_{k'}(x_1')=H'_{k'}(x_2')\wedge x_1'\neq x_2']\geq 1/p(n).$$

Moreover, the probability that $x_1'=x_1^*$ is $1/2^{2n}=\mathsf{negl}(n)$, so we have

$$\Pr_{k\leftarrow\{0,1\}^{m+2n},x_1^*,x_2^*\leftarrow\{0,1\}^{2n}}[A'(k,x_1^*,x_2^*,x_1')\to x_2' : H_{k}(x_1')=H'_{k'}(x_2')\wedge x_1'\neq x_2']\geq 1/p(n)-\mathsf{negl}(n).$$

Now, we have two cases, one of which must happen with at least $1/(2p(n))-\mathsf{negl}(n)$ probability: either (a) $x_2'=x_1^*$, or (b) not.

Suppose we are in case (a). Then, by definition of H', we have $H'_{k'}(x_2')=H'_{k'}(x_1^*)=H_k(x_2^*)$, and thus, with probability $1/(2p(n))-\mathsf{negl}(n)$, we have $H_k(x_1')=H_k(x_2^*)$ (and $x_1'\neq x_2'$). Since x_1' is fixed and x_2^* is uniformly random, this gives the following adversary A for $\mathcal{H}$ (for target x_1'): Simply output a uniformly random x_2^*. This will be distinct from x_1' with $1-1/2^{2n}$ probability, and will succeed with probability $1/(2p(n))-\mathsf{negl}(n)$. This contradicts universal one-wayness for $\mathcal{H}$, as desired.

Now, suppose we are in case (b). We can define $x_1:=x_1'$ and the following adversary A for $\mathcal{H}$: On input (k,x_1), uniformly randomly sample $x_1^*,x_2^*\leftarrow\{0,1\}^{2n}$, and run $A'(k,x_1^*,x_2^*,x_1')\to x_2'$ and output $x_2:=x_2'$. We know

$$H_k(x_1)=H'_{k'}(x_2)\wedge x_1\neq x_2$$

with probability $1/(2p(n))-\mathsf{negl}(n)$. Moreover, $x_2=x_2'\neq x_1^*$, so we know $H'_{k'}(x_2)=H_k(x_2)$. Therefore, x_1 and x_2 form a collision for $\mathcal{H}$, as desired.

Finally, we show that $\mathcal{H}'$ is not collision resistant. Consider the adversary that on input $k'=(k,x_1^*,x_2^*)$, outputs (x_1^*,x_2^*). By construction, $H'_{k'}(x_2^*)=H_k(x_2^*)=H'_{k'}(x_1^*)$, and with probability $1-1/2^{2n}$, $x_1^*\neq x_2^*$. Therefore, this adversary succeeds with probability $1-\mathsf{negl}(n)$, which is non-negligible.

(d) (2 points) Show that for all of these definitions, for all $k\in\{0,1\}^m$,

$$\Pr_{x\leftarrow\{0,1\}^{2n}}\left[\left|H_k^{-1}(H_k(x))\right|\leq 1\right]\leq\frac{1}{2^n}.$$

Solution

Clearly for all x, $\left|H_k^{-1}(H_k(x))\right|\geq 1$, so we only have to consider the case where $\left|H_k^{-1}(H_k(x))\right|=1$. Let $S\subseteq\{0,1\}^{2n}$ denote the set of all such x values where $\left|H_k^{-1}(H_k(x))\right|=1$.

Claim: $|S|\leq 2^n$. Suppose not, for the sake of contradiction. Then, since the codomain of H_k is $\{0,1\}^n$, by the pigeonhole principle, there must exist $x_1\neq x_2\in S$ such that $H_k(x_1)=H_k(x_2)$. Therefore, $\{x_1,x_2\}\subseteq H_k^{-1}(H_k(x_1))$, which is a contradiction with the fact that $|H_k^{-1}(H_k(x_1))|=1$.

Now, we have

$$\Pr_{x\leftarrow\{0,1\}^{2n}}\left[\left|H_k^{-1}(H_k(x))\right|\leq 1\right]=\Pr_{x\leftarrow\{0,1\}^{2n}}[x\in S]=\frac{|S|}{2^{2n}}\leq 1/2^n,$$

as desired.

(e) (3 points) Suppose $\mathcal{H}$ is a weak collision resistant hash function family. Show that the function $F : \{0,1\}^{m+2n} \to \{0,1\}^{m+n}$ given by $F(k, x) = (k, H_k(x))$ is a one-way function.

Solution

Clearly, F is efficiently computable, so it suffices to show that it is hard to find pre-images for F. Suppose for contradiction that there is some inverter A such that

$$\Pr_{(k_1,x_1)\leftarrow\{0,1\}^{m+2n}}[A(F(k_1, x_1)) \to (k_2, x_2) : F(k_1, x_1) = F(k_2, x_2)] \geq 1/p(n),$$

for some polynomial $p(n)$ infinitely often. Since $(k_1, H_{k_1}(x)) = F(k_1, x_1) = F(k_2, x_2) = (k_2, H_{k_2}(x_2))$, it is clear that if A succeeds, then since $k_1 = k_2 =: k$, we have $H_k(x_1) = H_k(x_2)$. That is,

$$\Pr_{(k,x_1)\leftarrow\{0,1\}^{m+2n}}[A(k, H_k(x_1)) \to (k, x_2) : H_k(x_1) = H_k(x_2)] \geq 1/p(n).$$

Now, consider the following adversary A' for weak collision resistance: on input (random) (k, x_1), run $A(k, H_k(x_1)) \to (k, x_2)$ and output x_2. Clearly, with probability $1/p(n)$, it holds that $H_k(x_1) = H_k(x_2)$, but we do not know whether $x_1 = x_2$.

From the previous part, we know that the probability over random x_2 that there exists $x_1 \neq x_2$ such that $H_k(x_1) = H_k(x_2)$ is $1 - \mathsf{negl}(n)$. Let X denote the set of pre-images of $H_k(x_2)$, i.e., making $|X| \geq 2$ with high probability. Note that for any element $x \in X$, the adversary behaves *identically*: it only sees k and $H_k(x)$, which is the same for all $x \in X$. Therefore, A's success probability is $1/p(n)$ *for all* $x \in X$. In particular, in the case where $|X| \geq 2$, the probability that $x_1 \neq x_2$ (conditioned on the success of A and $|X| \geq 2$) is at least $1/2$, making the overall success probability of the collision finder A' at least $1/(2p(n)) - \mathsf{negl}(n)$, which is non-negligible, as desired.

It turns out that the existence of one-way functions also implies the existence of weak collision resistant hash functions and universal one-way hash functions, even though it is an open question for collision resistant hash functions! You may assume these facts throughout the problem set.

Problem 2. Shrinking Signature Trees. (8 points)

Recall the many-time signature scheme using signature trees from class (see Lecture 12 slides). In the scheme presented in class, Alice computes many (VK, SK) pairs and arranges them in a tree of depth λ. In this problem, we will explore whether Alice can save space by computing only parts of the tree.

(a) (6 points) Suppose Alice signs messages using the leaves in order from left to right. Instead of generating the entire tree in advance, she proposes generating only the nodes along the path to the current leaf and their siblings, Show that Alice can use this approach to sign up to 2^λ messages while storing only $O(\lambda)$ signing keys, verification keys and signatures at any time. Briefly justify that the resulting signature scheme is EUF-CMA secure, assuming that the underlying one-time signature scheme is EUF-CMA secure.

Solution

Recall that in the signature tree construction, each internal node uses its signing key to sign the concatenation of the verification keys of its two children, and each leaf uses its signing key to sign a message.

Initially, Alice generates the nodes along the path from the root to the leftmost leaf and their siblings. Whenever she generates the two children of a node, she uses the signing key of the

node to sign their verification keys. She then uses the signing key of the leftmost leaf to sign the first message, as in the original construction.

Now suppose Alice has just signed a message using leaf i and wants to sign a message using leaf $i+1$. Let v be the lowest common ancestor of leaves i and $i+1$. Alice can discard all the keys and signatures in the left subtree of v except for the verification key of the left child of v, since all the leaves in this subtree have already been used. Also, Alice has already generated the signing and verification keys of the right child of v, and has a signature from v on the verification keys of its children. So starting from the right child of v, Alice can generate the remaining nodes along the path to leaf $i+1$ and their siblings, signing the verification keys of the newly generated children as before.

By repeating this process, Alice can sign 2^λ messages. At any point in time, Alice only needs to store the nodes along a root-to-leaf path and their siblings, of which there are at most $2\lambda + 1$. Since she stores at most a constant number of signing keys, verification keys, and signatures per node, the total number of each that she stores is $O(\lambda)$.

Finally, note that the only difference between this construction and the original construction is when the signing and verification keys are generated and discarded. Since the keys for each node are generated independently, the distribution of signatures that an adversary sees is identical to that in the original construction. Therefore, by the security of the original signature tree construction, the resulting scheme is EUF-CMA secure assuming that the underlying one-time signature scheme is EUF-CMA secure.

(b) (2 points) Now Alice wants to reduce her storage even further. She proposes generating the nodes along the path to the current leaf and their siblings as in part (a), but discarding the signing keys of the sibling nodes after their parent has signed their verification keys. When she later needs to use one of these sibling nodes, she generates a fresh signing and verification key pair for it. Explain why this approach does not work.

Solution

The problem is that when Alice later generates a fresh signing and verification key pair for a sibling node, the new verification key need not be the same as the original verification key. However, the signature from its parent is on the original verification key (together with that of its sibling), not the new one. Thus, Alice can no longer use this signature to authenticate the node.

Problem 3. Statistically Hiding Commitments (10 points)

Recall the commitment scheme based on an injective one-way function from class (Lecture 13 notes, Section 5).

(a) (3 points) Prove that no non-interactive commitment scheme without public parameters can be both statistically hiding and computationally binding.

Solution

Suppose there exists a non-interactive commitment scheme Com without public parameters that is both statistically hiding and computationally binding.

By statistical hiding, the statistical distance between commitments to 0 and commitments to 1 is negligible. In particular, the supports of the two commitment distributions must intersect

(otherwise their statistical distance would be 1). Thus, for every sufficiently large λ, there exist strings r_0 and r_1 such that

$$\mathsf{Com}(0, r_0) = \mathsf{Com}(1, r_1).$$

Recall that in class we considered security against non-uniform polynomial-size adversaries. Since there are no public parameters, we can define an adversary that, for each security parameter λ, hardcodes such a pair (r_0, r_1) and outputs it. This adversary has polynomial size since the length of the randomness is polynomial in λ. Moreover, the adversary succeeds in finding two openings of the same commitment to different bits with probability 1, contradicting computational binding, as desired.

(b) (7 points) Suppose that instead of an injective one-way function, we use a family of collision-resistant hash functions

$$H_{pp} : \{0,1\}^{2\lambda} \to \{0,1\}^{\lambda},$$

where pp is a publicly sampled hash-function key. Consider the following modification of the commitment scheme from class: to commit to $b \in \{0,1\}$, Alice samples independent uniform $r, s \in \{0,1\}^{2\lambda}$ and computes

$$\mathrm{Com}(pp, b; (r, s)) = (H_{pp}(r), s, \langle r, s\rangle \oplus b).$$

To open the commitment, Alice reveals r, and the receiver verifies the opening as in the original construction.

Prove that the resulting commitment scheme is statistically hiding and computationally binding.

Solution

First we show that the commitment scheme is computationally binding. Suppose there exists a PPT adversary A that, given pp, outputs a commitment $c = (y, s, t)$ and valid openings r_0 and r_1 to 0 and 1, respectively, with non-negligible probability.

By the verification procedure, if A succeeds, we have $H_{pp}(r_0) = y = H_{pp}(r_1)$ and $\langle r_0, s\rangle = t = \langle r_1, s\rangle \oplus 1$. In particular, the second equality implies that $r_0 \neq r_1$. Therefore, we can define an adversary A' for collision resistance as follows: on input pp, run $A(pp)$ and output (r_0, r_1). Clearly, A' is PPT if A is. Moreover, by the above argument, A' finds a collision for H_{pp} whenever A succeeds, which happens with non-negligible probability. This contradicts the collision resistance of H.

Now we show that the commitment scheme is statistically hiding. Fix any pp, and for each $y \in \{0,1\}^{\lambda}$, let $R_y := H_{pp}^{-1}(y)$. Note that, conditioned on $H_{pp}(r) = y$, r is uniformly distributed over R_y.

For each y such that $R_y \neq \emptyset$ and each $s \in \{0,1\}^{2\lambda}$, let

$$\epsilon_{y,s} := \Pr_{r \leftarrow R_y}[\langle r, s\rangle = 0] - \Pr_{r \leftarrow R_y}[\langle r, s\rangle = 1].$$

Equivalently, we have $\epsilon_{y,s} = \frac{1}{|R_y|}\sum_{r \in R_y}(-1)^{\langle r,s\rangle}$.

We first claim that $\mathbb{E}_{s \leftarrow \{0,1\}^{2\lambda}}[\epsilon_{y,s}^2] = \frac{1}{|R_y|}$. To see this, note that

$$\begin{aligned}\mathbb{E}_s[\epsilon_{y,s}^2] &= \frac{1}{|R_y|^2}\sum_{r,r' \in R_y}\mathbb{E}_s[(-1)^{\langle r \oplus r', s\rangle}] \\ &= \frac{1}{|R_y|}.\end{aligned}$$

The last equality follows from the fact that if $r = r'$, then the expectation is 1, and otherwise, $\langle r \oplus r', s\rangle$ is uniformly distributed over $\{0,1\}$, so the expectation is 0.

Therefore, by Cauchy–Schwarz, we have

$$\mathbb{E}_s[|\epsilon_{y,s}|] \le \sqrt{\mathbb{E}_s[\epsilon_{y,s}^2]} = \frac{1}{\sqrt{|R_y|}}.$$

Let D_b denote the distribution $(H_{pp}(r), s, \langle r, s\rangle \oplus b)$ for independent uniformly random $r, s \in \{0,1\}^{2\lambda}$. Note that the first two components are identically distributed in D_0 and D_1. Moreover, conditioned on these components being (y, s), the statistical distance between the distributions of the last component is exactly $|\epsilon_{y,s}|$.

Thus, denoting statistical distance by Δ, we have

$$\begin{aligned}
\Delta(D_0, D_1) &= \sum_{y: R_y \neq \emptyset} \mathop{\mathbf{Pr}}_r[H_{pp}(r) = y]\, \mathbb{E}_s[|\epsilon_{y,s}|] \\
&\le \sum_{y: R_y \neq \emptyset} \frac{|R_y|}{2^{2\lambda}} \cdot \frac{1}{\sqrt{|R_y|}} \\
&= \frac{1}{2^{2\lambda}} \sum_{y \in \{0,1\}^\lambda} \sqrt{|R_y|} \\
&\le \frac{1}{2^{2\lambda}} \sqrt{2^\lambda \sum_{y \in \{0,1\}^\lambda} |R_y|} \\
&= \frac{1}{2^{2\lambda}} \sqrt{2^\lambda \cdot 2^{2\lambda}} \\
&= 2^{-\lambda/2}.
\end{aligned}$$

Here the second inequality follows from Cauchy–Schwarz, and the second-to-last equality follows from the fact that the sets R_y partition $\{0,1\}^{2\lambda}$.

Since $2^{-\lambda/2} = \mathsf{negl}(\lambda)$ and the above holds for any pp, the commitment scheme is statistically hiding.

Problem 4. Hiding the Circuits in FHE (20 points)

Recall the definition of what it means for a fully homomorphic encryption (FHE) scheme to be compact and to be circuit-hiding. In this problem, we will think of FHE schemes capable of evaluating a class of Boolean circuits $\mathcal{C}$ computing functions f with k-bit inputs and a 1-bit output, for any $k = \mathsf{poly}(n)$.

Compactness. We say that a secret-key FHE scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ has compact ciphertexts if there is a polynomial p such that for any $C \in \mathcal{C}$, the bit-length of $\mathsf{Eval}(\mathsf{ek}, C, \mathsf{Enc}(\mathsf{sk}, x))$ is at most $p(n)$. That is, the size of the evaluated ciphertext depends neither on the input length $|x|$ nor on the circuit size $|C|$.

Function Hiding. We say that a secret-key FHE scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$ is function-hiding if for every PPT adversary A, any input x, and any pair of circuits C_0 and C_1 of the same size with $C_0(x) = C_1(x)$:

$$\begin{aligned}
\mathbf{Pr}[b' = b : (\mathsf{sk}, \mathsf{ek}) \leftarrow \mathsf{Gen}(1^n); \psi \leftarrow \mathsf{Enc}(\mathsf{sk}, x); b \leftarrow \{0,1\}; c_b \leftarrow \mathsf{Eval}(\mathsf{ek}, C_b, \psi); b' \leftarrow A(\mathsf{sk}, \mathsf{ek}, \psi, c_b)] \\
\le 1/2 + \mathsf{negl}(n)
\end{aligned}$$

(a) (1 points) Show that compactness does not necessarily imply function hiding. For this problem, you can assume the existence of an FHE scheme that satisfies compactness.

Solution

Given a compact FHE scheme $(\mathsf{Gen}, \mathsf{Enc}, \mathsf{Dec}, \mathsf{Eval})$, and the expressions of the functions f in $\{0,1\}^*$, one can construct a new compact FHE scheme $(\mathsf{Gen}', \mathsf{Enc}', \mathsf{Dec}', \mathsf{Eval}')$, where we take $\mathsf{Gen}' = \mathsf{Gen}$,$\mathsf{Enc}'(sk, m) = \mathsf{Enc}(sk, m)\|0^n$, i.e., we concatenate n zero bits to the original encryption. Accordingly, we define the $\mathsf{Dec}'(c\|b_n) = \mathsf{Dec}(c)$, i.e., we return the output from the original Dec on the ciphertext after removing the last n bits b_n. Finally, we define $\mathsf{Eval}'(f, c\|b^n) = \mathsf{Eval}(f, c)\|f_n$, where f_n is the first n bits of the expression of the function f. Thus, the size of evaluated ciphertext only increases by n bits, dependent on the security parameter only, and thus is still compact. However, the evaluated ciphertext leaks the information about the function f to be evaluated.

(b) (1 points) Show that function hiding does not necessarily imply compactness. For this problem, you can assume the existence of an FHE scheme that satisfies function hiding.

Solution

Similar to (a), it is noted that we can always artificially add dummy bits to the evaluated ciphtertext to make it not compact, while the decryption can simply ignore those dummy bits.

(c) (18 points) Assume the existence of a function hiding (but not necessarily compact) secret-key FHE scheme $(\mathsf{Gen}_1, \mathsf{Enc}_1, \mathsf{Dec}_1, \mathsf{Eval}_1)$ and that of a compact (but not necessarily function-hiding) secret-key FHE scheme $(\mathsf{Gen}_2, \mathsf{Enc}_2, \mathsf{Dec}_2, \mathsf{Eval}_2)$. Construct a secret-key FHE scheme that is (IND-CPA secure) and both compact and function hiding.

For the purposes of this problem, we only require that the decryption algorithm decrypts evaluated ciphertexts (as opposed to fresh ciphertexts produced by the encryption algorithm) correctly.

Solution

We consider the following construction. Given a function f and a plaintext m,

- First, we use the Gen_1 and Gen_2 to generate (sk_1, ek_1) and (sk_2, ek_2), respectively.
- Second, for plaintext m, we use the function-hiding encryption Enc_1 to obtain the ciphtertext $c = \mathsf{Enc}_1(sk_1, m)$. Simultaneously, we obtain $\psi = \mathsf{Enc}_2(sk_2, sk_1)$, the compact encryption of the function-hiding encryption key sk_1. It is not hard to prove that the encryption on both m and ψ satisfies IND-CPA, as we simply use the encryption Enc_1 on m while use the independent Enc_2 on the secret key sk_1 of Enc_1.
- Third, we define the evaluation function as $\mathsf{Eval}(f, c)$ with the function f on ciphertext c by the following two steps: a). we compute $C_1 = \mathsf{Eval}_1(ek_1, f, c)$; b). we consider the decryption function $\mathsf{Dec}_1(\cdot, C_1)$, as a function of the secret key sk_1, and define $C_2 = \mathsf{Eval}_2(\mathsf{ek}_2, \mathsf{Dec}_1(\cdot, C_1), \psi)$, to evaluate the circuit of $\mathsf{Dec}_1(\cdot, C_1)$ under the compact FHE scheme. C_2 is our final evaluation ciphertext. This satisfies compactness requirement, since C_2 is produced by Eval_2, and thus the size C_2 is some polynomial only dependent on the security parameter n by our compactness assumption of Eval_2, but independent to the input from Enc_1 and the function $\mathsf{Dec}_1(\cdot, C_1)$. Additionally, from the function-hiding assumption of Enc_1, our evaluation function is a function of the ciphertext from Enc_1, and thus we satisfy the function-hiding property.
- Finally, we define the decryption as $\mathsf{Dec}_2(sk_2, C_2)$, which returns m. To see the correctness, it is noted that $C_2 = \mathsf{Eval}_2(\mathsf{ek}_2, \mathsf{Dec}_1(\cdot, C_1), \psi)$ actually equals $\mathsf{Enc}_2(sk_2, f(m))$.