Final

Name:

SID:

- Not all parts of a problem are weighted equally. Before you answer any question, read the problem carefully. Be precise and concise in your answers.
- You may consult at most *10 sheets of notes*. Apart from that, you may not look at books, notes, etc. Calculators, phones, computers, and other electronic devices are NOT permitted for looking up content. However, you may use an electronic device such as a tablet for writing your answers.
- You have 170 minutes: there are 5 questions on this exam worth a total of 100 points.
- You are allocated 180 minutes and the extra 10 minutes are provided for the submission of the exam on Gradescope. You must submit/upload the exam on time. Note that late submissions will not be accepted.
- **DSP Students** must submit the exam in time as per your accommodation. Thus, if you are allowed 1.5× (resp., 2×) the exam time then you must submit it within $170 * 1.5 + 10$ (resp., $170 * 2 + 10$) mins, i.e. 265 (resp., 350) mins. **Your allotted time has already been adjusted on Gradescope, and you will not need to email the final to the instructors. Please submit on Gradescope.**
- The exam must be submitted before 8 AM PT on May 12th, 2021. No exams later than this time will be accepted regardless of when you start the test.
- We will not be answering questions during the exam. If you feel that something is unclear please write a note in your answer.

1 Fill in the blanks (10 Points)

Please write the term that each sentence is describing. Each is worth 1 point.

(1) A function from natural numbers to non-negative real numbers that is asymptotically smaller than $1/x^c$ for any positive integer c.

Solution: negligible function

(2) A method of obtaining public-key encryption at the asymptotic cost of private-key encryption.

Solution: hybrid encryption

(3) An attack model where the adversary can obtain decryptions of ciphertexts of its own choice.

Solution: chosen-ciphertext attack (CCA) model

(4) A two-party protocol to generate a shared secret key that can't be guessed by an eavesdropper.

Solution: secure key-exchange protocol

(5) An encryption scheme that allows for the following: given a ciphertext c encrypting a message m and a function f, compute a ciphertext encrypting $f(m)$.

Solution: homomorphic encryption scheme/fully homomorphic encryption scheme

(6) A proof technique used to show that two distributions are computationally indistinguishable (i.e., $\mathcal{D}_0 \overset{c}{\approx} \mathcal{D}_1$) by constructing a sequence of polynomially many distributions $\mathcal{H}_0, \mathcal{H}_1, \ldots, \mathcal{H}_n$ where $\mathcal{H}_0 = \mathcal{D}_0$ and $\mathcal{H}_n = \mathcal{D}_1$ and showing $\mathcal{H}_{i-1} \overset{c}{\approx} \mathcal{H}_i$ for all $i \in [n]$.

Solution: hybrid argument

(7) A bijective function that can be computed efficiently on every input, but cannot be inverted in polynomial time given the image of a random input.

Solution: one-way permutation

(8) DES and AES are examples of

Solution: block cipher

(9) For a cyclic group $\mathcal{G}$ of order q and with generator g, the assumption that it is computationally hard to find g^{ab} given (g, g^a, g^b), where $a, b \xleftarrow{\$} \mathbb{Z}_q$.

Solution: computational Diffie-Hellman (CDH) assumption

(10) An example of a perfectly secure encryption scheme.

Solution: one-time pad

Name:

2 True/False (20 points)

Bubble in the right answer. No explanation needed. +2 points for correct answer and **0 points for wrong answers/unanswered questions**! This part will be graded automatically. Please mark your answer clearly.

1. It is typically a good idea to hide the implementation details of a public-key encryption scheme. **Solution:** False

   ◯ True

   ◯ False

2. We can construct a perfectly secure encryption scheme with messages that are double the length of keys; e.g. $|\mathcal{M}| = |\mathcal{K}|^2$, where $|\mathcal{K}|$ and $|\mathcal{M}|$ are the key and message spaces respectively. **Solution:** False

   ◯ True

   ◯ False

3. If f is PRF, then $G(s) = s||f_s(1)||\ldots||f_s(n)$ (for $n > 0$) must be a PRG. **Solution:** False

   ◯ True

   ◯ False

4. If f, g are PRGs then $g(f(\cdot))$ must be a PRG. **Solution:** True

   ◯ True

   ◯ False

5. The "authenticate then encrypt" paradigm results in a secure authenticated encryption scheme. **Solution:** False

   ◯ True

   ◯ False

6. CPA-secure public-key encryption for one-bit messages implies CPA-secure public-key encryption for n-bit messages (for $n \geq 2$). **Solution:** True

   ◯ True

   ◯ False

7. We can construct a collision resistant hash function assuming the hardness of the LWE problem. **Solution:** True

   ◯ True

   ◯ False

8. It is reasonable to assume that the DDH problem is hard in groups with an efficient pairing function. **Solution:** False

   ◯ True

   ◯ False

9. Identity-based encryption implies CCA2 secure public-key encryption. **Solution:** True

   ◯ True

   ◯ False

10. A deterministic unforgeable MAC is always strongly unforgeable. **Solution:** True

    ◯ True

    ◯ False

Name:

3 One-Way Functions (20 points)

1. Show that the existence of a *non-interactive perfectly binding bit commitment scheme* implies a one-way function (10 points).

   Recall that a non-interactive perfectly binding bit commitment scheme Com is an algorithm that takes in a bit $b \in \{0,1\}$ and random coins $r \in \{0,1\}^n$, and produces a commitment c. It satisfies the following two properties.

   - **Hiding**, which states that no PPT adversary can distinguish between $\mathsf{Com}(0;r)$ and $\mathsf{Com}(1;r)$ for uniformly sampled $r \leftarrow \{0,1\}^n$.
   - **Perfect binding**, which states that there do not exist any r_0, r_1 such that $\mathsf{Com}(0;r_0) = \mathsf{Com}(1;r_1)$.

   Solution: Let $f(x) = \mathsf{Com}(x_1; x_{2,\ldots,n})$. Assume an adversary $\mathcal{A}$ that can invert $f(x)$ with non-negligible probability. By perfect binding, the first bit of the pre-image x' that $\mathcal{A}$ finds must be equal to the first bit of x. However, this immediately implies that $\mathcal{A}$ can break the hiding of Com, by inverting the commitment c received with inverse polynomial probability.

2. Let $(\mathsf{Gen}, \mathsf{Sign}, \mathsf{Vrfy})$ be a *perfectly correct* secure digital signature scheme. Perfect correctness states that for any message m,

$$\Pr_{r_{\mathsf{Gen}}, r_{\mathsf{Sign}} \leftarrow \{0,1\}^n, (\mathsf{vk},\mathsf{sk}) := \mathsf{Gen}(1^n; r_{\mathsf{Gen}})} [\mathsf{Vrfy}(\mathsf{vk}, m, \mathsf{Sign}(\mathsf{sk}, m; r_{\mathsf{Sign}})) = 1] = 1,$$

   where r_{Gen} are the random coins used by Gen and r_{Sign} are the random coins used by Sign.

   (a) Define $f(x)$ to output the verification key vk output by $\mathsf{Gen}(1^n; x)$. Show that f is a one-way function (5 points). **Solution:**

   (a) If there exists a PPT $\mathcal{A}$ that can invert f with non-negligible probability, then we can construct a PPT $\mathcal{B}$ that breaks the security of the signature scheme:
   (i) $\mathcal{B}$ gets pk from its challenger and forwards it to $\mathcal{A}$.
   (ii) $\mathcal{A}$ outputs x' such that $f(x') = \mathsf{pk}$.
   (iii) $\mathcal{B}$ computes $(\mathsf{pk}, \mathsf{sk}') := \mathsf{Gen}(1^n; x')$.
   (iv) $\mathcal{B}$ picks an arbitrary message m and computes $\sigma \leftarrow \mathsf{Sign}_{\mathsf{sk}'}(m)$. Output (m, σ).
   Since $(\mathsf{pk}, \mathsf{sk}')$ is generated from Gen, σ is a valid signature for m with respect to pk. Hence $\mathcal{B}$ breaks the security of the signature scheme with non-negligible probability.

   (b) Show that there exists a secure digital signature scheme $(\mathsf{Gen}', \mathsf{Sign}', \mathsf{Vrfy}')$ that is *not* perfectly correct (i.e. it is only correct with probability $1 - \mathsf{negl}(n)$), for which f as defined above is *not* a one-way function (5 points). **Solution:** $\mathsf{Gen}'(1^n; x)$ parses $x = (x_1, x_2) \in \{0,1\}^{2n}$ and outputs x_2 if $x_1 = 0^n$, and otherwise runs $\mathsf{Gen}(1^n; x_2)$. Note that Gen' will only fail to run Gen with negligible probability, but that any image y of $f(\cdot)$ has pre-image $(0, y)$.

4 Witness Indistinguishability (10 points)

Let L be a language in NP (e.g. graph three coloring) and let R_L be the NP-relation defined by the language L, meaning that $x \in L$ iff there exists a witness w (e.g. the graph three coloring function) such that $(x, w) \in R_L$. Let (P, V) be an interactive proof system for L. That is, both P and V are initialized with an instance x, the prover P is additionally given a witness w such that $(x, w) \in R_L$, and P attempts to convince V that $x \in L$. They do so by interacting, and we let $\langle P(w), V\rangle(x)$ denote the verifier's *view* of this interaction, consisting of the messages sent back and forth, as well as the verifier's private state.

We say that (P, V) is *witness indistinguishable (WI)* if for all adversarial PPT V^*, all $x \in L$, and all distinct witnesses w_0, w_1 such that $(x, w_0) \in R_L$ and $(x, w_1) \in R_L$, the following two distributions are computationally indistinguishable.

$$\langle P(w_0), V^*\rangle(x) \approx \langle P(w_1), V^*\rangle(x).$$

1. Show that if (P, V) is a computational *zero-knowledge* proof system, then it is also witness indistinguishable (5 points).

   Solution: The zero-knowledge property implies that there exists a simulator Sim such that for any witnesses w, it holds that $\langle P(w), V^*\rangle(x) \approx \mathsf{Sim}(x)$. This in particular means that $\langle P(w_0), V^*\rangle(x) \approx \mathsf{Sim}(x)$ and that $\langle P(w_1), V^*\rangle(x) \approx \mathsf{Sim}(x)$, and the claim follows.

2. Let (P, V) be a witness-indistinguishable proof system. Define $(\widetilde{P}, \widetilde{V})$ to repeat (P, V) independently k times *in parallel* (where k is some polynomial), and $\widetilde{V}$ accepts if and only if V accepts in all the parallel executions. Show that $(\widetilde{P}, \widetilde{V})$ is still witness indistinguishable (5 points). **Solution:** We will define a sequence of hybrids where Hyb_i uses the witness w_1 in the first i parallel repetitions and uses the witness w_0 in the rest. Notice that Hyb_0 corresponds to the distribution where w_0 is used in every repetition and Hyb_k corresponds to the distribution where w_1 is used in every repetition. It follows from witness indistinguishability that Hyb_i is computationally indistinguishable to Hyb_{i+1} for all i and the claim follows.

Name:

5 CCA-Secure Encryption (40 points)

An *injective trapdoor function* f is a keyed injective one-way function that can be inverted given a "trapdoor" td associated with the public key k. It consists of the following three algorithms.

- $\mathsf{Gen}(1^n) \to (k, \mathsf{td})$ outputs a key and trapdoor pair.
- $f_k(x) = y$ evaluates the function on input x.
- $f_k^{-1}(\mathsf{td}, y) = x$ inverts the function on output y, given the trapdoor td.

Correctness states that for any $(k, \mathsf{td}) \leftarrow \mathsf{Gen}(1^n)$ and $x \in \{0,1\}^n$, it holds that $f_k^{-1}(\mathsf{td}, f_k(x)) = x$. One-wayness states that for any PPT $\mathcal{A}$,

$$\Pr_{(k,\mathsf{td})\leftarrow\mathsf{Gen}(1^n), x\leftarrow\{0,1\}^n}[\mathcal{A}(k, f_k(x)) \to x] = \mathsf{negl}(n).$$

We now define a stronger security property for injective trapdoor functions, which we call *correlated input security*. Intuitively, this states that the adversary, given $f_{k_i}(x)$ for each $i \in \{1...n\}$ (where each k_i is sampled independently), cannot recover x. Formally, for any PPT $\mathcal{A}$,

$$\Pr_{\{(k_i,\mathsf{td}_i)\leftarrow\mathsf{Gen}(1^n)\}_{i\in[n]}, x\leftarrow\{0,1\}^n}[\mathcal{A}(k_1, \ldots, k_n, f_{k_1}(x), \ldots, f_{k_n}(x)) \to x] = \mathsf{negl}(n).$$

Let $h(\cdot)$ be a hard-core predicate for a correlated input secure injective trapdoor function f. Now consider the following public-key encryption scheme.

- $\mathsf{PKE.Gen}(1^n)$ samples $2n$ keys $\{(k_i^0, \mathsf{td}_i^0) \leftarrow \mathsf{Gen}(1^n), (k_i^1, \mathsf{td}_i^1) \leftarrow \mathsf{Gen}(1^n)\}_{i\in[n]}$ and defines the secret key and public key as

$$\mathsf{pk} := ((k_1^0, k_1^1), \ldots, (k_n^0, k_n^1))$$
$$\mathsf{sk} := ((\mathsf{td}_1^0, \mathsf{td}_1^1), \ldots, (\mathsf{td}_n^0, \mathsf{td}_n^1))$$

- $\mathsf{PKE.Enc}(\mathsf{pk}, m)$ for $m \in \{0,1\}$ samples $v \leftarrow \{0,1\}^n$ and $x \leftarrow \{0,1\}^n$ and outputs $\mathsf{ct} = (v, y_1, \ldots, y_n, c)$, where each $y_i = f_{k_i^{v_i}}(x)$, and $c = h(x) \oplus m$.

Now answer the following questions. Note that an answer to part 3 is not necessarily required to answer parts 4 and 5.

1. Prove that correlated input security implies one-wayness (5 points). **Solution:** An adversary that can break one-wayness can be used to break correlated input security as follows. The reduction will forward one of its images (e.g. $f_{k_1}(x)$) to the one-wayness adversary, who returns x with inverse polynomial probability. Since f is injective, this x would be the pre-image for the reduction's challenge.

2. Give a description of $\mathsf{PKE.Dec}(\mathsf{sk}, \mathsf{ct})$ to go along with $\mathsf{PKE.Gen}, \mathsf{PKE.Enc}$ defined above (5 points). **Solution:** $\mathsf{PKE.Dec}(\mathsf{sk}, \mathsf{ct})$ inverts each y_i with $\mathsf{td}_i^{v_i}$ to obtain x_i. If $x_1 = \cdots = x_n = x$, output $m = c \oplus h(x)$.

3. Show that, assuming f is a correlated input secure injective trapdoor function with hard-core predicate $h(\cdot)$, this scheme satisfies CCA1 security (recall that CCA1 security does not give the adversary access to the decryption oracle after the challenge phase) (10 points). **Solution:** Consider the two games where during the challenge phase, the adversary's query $m \in \{0,1\}$ is either answered as the honest encryption of m (Game 1) or as the encryption of a random bit (Game 2). The reduction is given keys $k_1, \ldots, k_n$, an output $y_1 = f_{k_1}(x), \ldots, y_n = f_{k_n}(x)$, and a bit b that is either $h(x)$ or a random bit. It samples $v^* \leftarrow \{0,1\}^n$, sets $k_i^{v_i^*} = k_i$ for each $i \in [n]$, and samples fresh (key, trapdoor) pairs for each $k_i^{1-v_i^*}$. Then it sends $((k_1^0, k_1^1), \ldots, (k_n^0, k_n^1))$ as the public key of the PKE scheme to the CCA1 adversary. Note that v^* is independent of the adversary's view, and that the reduction can succesfully answer decryption queries $(v, y_1, \ldots, y_n, c)$ as long as $v \neq v^*$, which occurs with all but $\mathsf{negl}(n)$ probability. During the challenge phase, the reduction receives m and encrypts honestly, except that b is used in place of $h(x)$. Note that if b was the hard-core bit, then we are in Game 1, and if b was random, then we are in Game 2. But distinguishing these games implies the reduction's ability to distinguish the hard-core bit $h(\cdot)$ from random, a contradiction.

4. Give the formal security definition of a one-time strongly unforgeable signature scheme $(\mathsf{Sig.Gen}, \mathsf{Sig.Sign}, \mathsf{Sig.Vrfy})$ (recall that "one-time" security states that security holds as long as the adversary only sees a signature on a single message, and "strong unforgeablity" states that the adversary cannot even produce a different signature on a message that it queried to its signing oracle) (5 points). **Solution:** The adversary $\mathcal{A}$ is given the public key pk and is able to query the signing oracle on exactly one message m, receiving $\sigma \leftarrow \mathsf{Sign}(\mathsf{sk}, m)$. Then they output (m', σ') and win if $\mathsf{Vrfy}(\mathsf{pk}, m', \sigma') = 1$ and $(m', \sigma') \neq (m, \sigma)$. Security states that no PPT $\mathcal{A}$ can win except with $\mathsf{negl}(n)$ probability.

5. Show how to use a one-time strongly unforgeable signature scheme to modify the public-key encryption scheme given above so that it achieves CCA2 security (where the adversary is additionally given decryption oracle access after the challenge phase) (5 points). **Solution:** Rather than sampling $v \leftarrow \{0,1\}^n$, sample $(\mathsf{vk}, \mathsf{sk}) \leftarrow \mathsf{Sig.Gen}(1^n)$ and set $v = \mathsf{vk}$. Then the ciphertext is $(v, y_1, \ldots, y_n, c, \sigma)$, where $\sigma \leftarrow \mathsf{Sig.Sign}(\mathsf{sk}, (y_1, \ldots, y_n, c))$. During decryption, first verify that σ is a valid signature on $(y_1, \ldots, y_n, c)$ under key v, and if not output $\perp$.

6. Prove that your modification of the scheme satisfies CCA2 security (10 points). **Solution:** Let the challenge ciphertext returned to the adversary be $(v^*, y_1^*, \ldots, y_n^*, c^*, \sigma^*)$. Note that as long as $v \neq v^*$ in each of the adversary's phase 2 queries, the argument from part 2 suffices to show CCA2 security. Thus, it remains to handle queries that begin with v^*. For such queries, we know that the remainder of the ciphertext $(y_1, \ldots, y_n, c, \sigma) \neq (y_1^*, \ldots, y_n^*, c^*, \sigma^*)$. Thus, if

σ^* is a valid signature on the remainder of the ciphertext under key v^*, the adversary can be used to break the security of the one-time strongly unforgeable signature scheme. Thus, each of these queries can be answered with $\perp$.