Final Exam

1 Problem 1 (100 points)

Answer the following. We expect your answer to each part to be roughly 1 to 3 sentences.

(a) For many message authentication codes, including some we saw in class, verification works as follows: re-compute the correct signature, and check for equality between the claimed signature and correct signature. However, none of the signature schemes we've seen work this way. Explain why.

(b) A friend claims to have designed a super-secure block cipher, explaining that it must be far more secure than AES because it has million-bit keys. Explain what is wrong with your friend's reasoning.

(c) Consider the following encryption scheme with key space $\{0,1\}^n$, message space $\{0,1\}$, and ciphertext space $\{0,1\}^n$. To encrypt a bit m with key k, choose a random $c \in \{0,1\}^n$ such that $c \cdot k = m$, where $c \cdot k = \sum_{i=1}^{n} c_i k_i \bmod 2$. Does the scheme have perfect secrecy (for a single message)? Explain.

(d) Any PRF with a sufficiently large range is also a MAC. Give an example of a MAC that is not a PRF, and explain why. You do not need to formally prove the security of your MAC.

(e) 3DES fixes the key length problem of DES, but 3DES still has a small block size. What is wrong with the following attempt at increasing the block size: set $E((k_0,k_1),(x_0,x_1)) = (3DES(k_0,x_0), 3DES(k_1,x_1))$? That is, the new block cipher will split the input into two parts, and apply 3DES to each part (each part getting independent keys).

(f) Explain why it is not possible to have a secure commitment scheme with 16-bit commitments, even if you only want to commit to single bit messages.

(g) The Schnorr proof of knowledge for discrete log is zero knowledge against honest verifiers. Suppose that the Schnorr proof of knowledge was additionally zero knowledge for *malicious* verifiers. Briefly explain why this would imply that Schnorr signatures are insecure. Specifically, explain how a malicious verifier, combined with the supposed simulator, can be used to forge Schnorr signatures.

(h) In Project 3, you explored how to attack ElGamal by looking at the time it takes to decrypt a handful of carefully crafted ciphertexts. The attack leverages timing variations in modular exponentiation. Modular exponentiation is also used during encryption. Explain why such timing variations during encryption are much less concerning.

(i) Suppose you have a secure signature scheme that works on 64-bit messages, and you want to extend the message space to larger messages by first hashing with a collision-resistant hash function and then signing just the hash of the message. Explain what this approach will not work.

(j) Your friend claims, for t-out-of-n secret sharing, that the prime p must be exponentially large, or else the attacker can brute-force search for everyone's shares and therefore recover the secret. Explain why this reasoning is flawed.

2 Problem 2 (50 points)

Answer the following. We expect your answer to each part to be roughly 1 to 3 sentences. Remember that if a number e is relatively prime to $\phi(N)$, then the map $x \mapsto x^e \bmod N$ is a permutation on $\mathbb{Z}_N^*$.

(a) Suppose it is discovered tomorrow that factoring integers is easy, but solving discrete log mod primes remains hard. What can you say about solving discrete log mod N, where N is the product of two large primes?

(b) Consider the equation $x^3 - 6x^2 + 12x - 6 = 0 \bmod N$, where N is the product of two large primes. Suppose 3 is relatively prime to $\phi(N)$. How many solutions does the equation have mod N? Explain.

(c) Consider the equation $x^3 - 6x^2 + 11x - 6 = 0 \bmod N$, where N is the product of two large primes. Suppose 3 is relatively prime to $\phi(N)$. How many solutions does the equation have mod N? Explain.

(d) Suppose N is the product of two unknown large primes, and that 15 is relatively prime to $\phi(N)$. Explain how to sample a random $x \in \mathbb{Z}_N^*$, together with a $y \in \mathbb{Z}_N^*$ such that $x^3 = y^5 \bmod N$.

(e) Let p be an odd prime, and suppose -1 is *not* a quadratic residue mod p (which is equivalent to $p = 3 \bmod 4$). Explain how to tell if a number x is a *quartic* residue mod p. We say that x is a quartic residue mod p if there exists an integer y such that $y^4 = x \bmod p$.

3 Problem 3 (50 points)

Recall the discrete log problem: given a group $\mathbb{G}$ of prime order p, and random $g, h \in \mathbb{G}$ such that $g \neq 1$ (meaning g is a generator), compute $a \in \mathbb{Z}_p$ such that $h = g^a$.

Now consider the *n-Vector Discrete Log* (n-VDL) problem: given a group $\mathbb{G}$ of prime order p and n random group elements $(h_1, h_2, \ldots, h_n)$, compute any vector $\mathbf{v} = (v_1, \ldots, v_n) \in \mathbb{Z}_p^n$ such that (1) $h_1^{v_1} \cdot h_2^{v_2} \cdot \cdots \cdot h_n^{v_n} = 1 \in \mathbb{G}$, and (2) $(v_1, \ldots, v_n) \neq (0, \ldots, 0)$.

In this problem, you will show that the discrete log problem in $\mathbb{G}$ and the n-VDL problem in $\mathbb{G}$ are essentially equivalent.

(a) Suppose you have a discrete log adversary A, which succeeds with non-negligible probability ϵ. Construct an n-VDL adversary B. B should run A exactly once, and should succeed with probability at least $\epsilon - (1/p)$.

(b) Now suppose you have an n-VDL adversary B with non-negligible success probability ϵ. Consider the following discrete log adversary $A_i(g, h)$: Choose random $r_1, \ldots, r_n$ and set $h_i = h \times g^{r_i}$, and $h_j = g^{r_j}$ for $j \neq i$. Then run $\mathbf{v} \leftarrow B(h_1, \ldots, h_n)$.

Finish the description of A_i: supposing $\mathbf{v}$ is a valid solution to the n-VDL instance $(h_1, \ldots, h_n)$, explain how A_i will attempt to use $\mathbf{v}$ to compute the discrete log a such that $h = g^a$. Under what conditions on $\mathbf{v}$ will A_i succeed?

(c) Fix some i. Explain why, even though B has non-negligible success probability ϵ, it may be possible that A_i fails to be a discrete log adversary (that is, that A_i may end up with a negligible or even zero probability of solving the discrete logarithm).

(d) While any particular A_i may fail, show that at least one of the A_i will have probability at least ϵ/n of solving the discrete logarithm. Thus at least one of the A_i will be a discrete log adversary.

(e) Devise a new adversary A which is guaranteed to succeed with probability at least $\epsilon(1 - 1/p)$. A should run B exactly once. To do so, your A will need to generate the $h_1, \ldots, h_n$ to feed to B in a way that guarantees A can compute a from $\mathbf{v}$ with overwhelming probability.

4 Problem 4 (60 points)

Consider the following public key identification scheme. Alice has the secret key, which consists of two large random primes p, q; Bob has the public verification key,

which consists of $N = pq$. To authenticate, Bob will send a random quadratic residue $x \in \mathbb{Z}^*$ to Alice. Alice will then respond with a square root y of x. Bob will accept if and only if $y^2 = x \bmod N$.

(a) Explain how Bob chooses the random quadratic residue x.

(b) The above description does not specify how Alice chooses y from among the 4 possible square roots of x. Show that, no matter how Alice chooses y, the protocol is *insecure* against *active* attacks.

(c) Show that the protocol is secure against direct attacks, under the assumption that factoring N is hard. Security should follow, no matter how Alice chooses y.

(d) Suppose Bob is legally required to record all incoming and outgoing messages, in case law enforcement needs a record for an investigation. Alice would like to be able to deny that any identification took place. To do so, she wants to enable Bob to generate, all by himself and *without* interacting with Alice, a transcript that looks exactly like the sequence of messages that would have been exchanged if Alice *had* identified herself to Bob. Bob should be able to do this, despite only knowing N and not p, q. Thus, law enforcement would be unable to tell if Alice actually identified herself, or if Bob simply generated the transcript of the identification by himself.

How should Alice choose the square root y in order to guarantee this property? Explain how Bob samples random-looking transcripts, and prove that Bob's sampled transcripts are distributed identically to actual transcripts of interaction with Alice.

(e) Prove that the identification protocol is secure against eavesdropping attacks, using Alice's method for choosing y from part (d).

5 Problem 5 (40 points)

Let $F : \{0,1\}^n \to \{0,1\}^m$ be a secure one-way function, where n, m are functions of the security parameter λ. For an integer $i \geq 0$, let $F_{\leq i}(x) = (F(x), x_1, \ldots, x_i)$, where x_j is the jth bit of x.

(a) Suppose $F_{\leq i}$ is a secure one-way function for some integer i. Prove that $F_{\leq (i+1)}$ is also a secure one-way function.

(b) Show that $F_{\leq n}$ is *not* one-way.

(c) Your friend thinks you made a mistake: "$F_{\leq 0}$ is just F, which is assumed to be secure. Supposing (a) is true, we can apply (a) n times, for $i = 0, 1, \ldots, n-1$. The result is that $F_{\leq n}$ is a secure one-way function, contradicting (b)."

Explain what is wrong with your friend's reasoning.