Problem Set 7 Solutions

Problem 20. *Do the following without a calculator or computer, showing your work. None should be overly tedious.* ***(a)*** *Compute* $7^{890002} \bmod 1111$. *Note that* 1111 *is the product of primes* $p = 11$ *and* $q = 101$. ***(b)*** *Compute* $7^{890002} \bmod 101$. ***(c)*** *Compute* $2^{64} \bmod 101$.

The answers are **49**, **49**, and **79**.

- We can take the exponent modulo the size of the group $\mathbb{Z}_N^*$, which is $\phi(N) = (p-1)(q-1) = 1000$. So the quantity we're interested in is $7^2 \bmod 1111$, which is 49.
- We can take the exponent modulo the size of the group $\mathbb{Z}_{101}^*$, which is 100, it's again 49.
- We "power-up." All that follows is in $\mathbb{Z}_{101}^*$:

$$\begin{aligned} 2^4 &= 16 \\ 2^8 &= 256 = 54 \\ 2^{16} &= 2^8 \cdot 2^8 = 54 \cdot 54 = 88 \\ 2^{32} &= 2^{16} \cdot 2^{16} = 88 \cdot 88 = 68 \\ 2^{64} &= 2^{32} \cdot 2^{32} = 68 \cdot 68 = 79 \end{aligned}$$

Problem 21. *Consider the following "left-or-right" security notion for a public-key encryption scheme* $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$*:*

$$\mathbf{Adv}^{\mathrm{lr}}_{\Pi}(A,k) \;=\; \Pr[(pk,sk) \twoheadleftarrow \mathcal{K}(k) \colon A^{\mathcal{E}(pk,\mathrm{L}(\cdot,\cdot))}(pk) \Rightarrow 1] - \Pr[(pk,sk) \twoheadleftarrow \mathcal{K}(k) \colon A^{\mathcal{E}(pk,\mathrm{R}(\cdot,\cdot))}(pk) \Rightarrow 1]$$

where oracle $\mathrm{L}(X,Y)$ *returns* X *when* $|X| = |Y|$*; oracle* $\mathrm{R}(X,Y)$ *returns* Y *when* $|X| = |Y|$*; and both oracles return* $\perp$ *when* $|X| \neq |Y|$*. In contrast, our old security notion was*

$$\mathbf{Adv}^{\mathrm{ind}}_{\Pi}(A,k) \quad = \quad \Pr[(pk,sk) \twoheadleftarrow \mathcal{K}(k) \colon A^{\mathcal{E}(pk,\cdot)}(pk) \Rightarrow 1] - \Pr[(pk,sk) \twoheadleftarrow \mathcal{K}(k) \colon A^{\mathcal{E}(pk,0^{|\cdot|})}(pk) \Rightarrow 1]$$

Show that lr-security is equivalent to ind-security.

We first show lr-security implies ind-security. As usual, we show such a statement by showing the contrapositive: if Π is insecure in the ind-sense then it is insecure in the lr-sense. Insecurity in the ind-sense means that there is an adversary A that gets good advantage for the ind definition. From A, we construct an adversary B that attacks Π in the lr-sense. Adversary B takes as input a public key pk and works as follows:

1. Run A. When A makes a query M, adversary B queries its own oracle with $(M, 0^{|M|})$ and forwards this reply to A.
2. When A output a bit and halts, adversary B outputs the same bit and halts.

If B was interacting with its left oracle, then it would simulate a real world oracle for A. If B was interacting with its right oracle, then it would simulate the ideal oracle for A. Formally, we have the following:

$$\Pr[B^{\mathcal{E}(pk,\text{Left}(\cdot,\cdot))}(pk) \Rightarrow 1] = \Pr[A^{\mathcal{E}(pk,\cdot)}(pk) \Rightarrow 1]$$
$$\Pr[B^{\mathcal{E}(pk,\text{Right}(\cdot,\cdot))}(pk) \Rightarrow 1] = \Pr[A^{\mathcal{E}(pk,0^{|\cdot|})}(pk) \Rightarrow 1]$$

which leaves us with

$$\begin{aligned}\mathbf{Adv}^{\text{lr}}_{\Pi}(A,k) &= \Pr[B^{\mathcal{E}(pk,\text{Left}(\cdot,\cdot))}(pk) \Rightarrow 1] - \Pr[B^{\mathcal{E}(pk,\text{Right}(\cdot,\cdot))}(pk) \Rightarrow 1] \\ &= \Pr[A^{\mathcal{E}(pk,\cdot)}(pk) \Rightarrow 1] - \Pr[A^{\mathcal{E}(pk,0^{|\cdot|})}(pk) \Rightarrow 1] \\ &= \mathbf{Adv}^{\text{ind}}_{\Pi}(A,k)\end{aligned}$$

Now for this direction, constructing an ind-adversary A from an lr-adversary B. Adversary A takes as input the public key pk and does the following:

1. $b \twoheadleftarrow \{0,1\}$
2. Run B. When B makes a query (M_1, M_0), query the ind-oracle with M_b. Respond to B with the oracle's response. Let B's output be the bit d.
3. If $b = d$, then return 1; otherwise, return 0.

When A is interacting with its ideal oracle, then the responses that A forwards to B are independent of B's queries. This means that B's chances of guessing b is only as good as randomly guessing, which is $1/2$. On the other hand, if A is interacting with its real oracle, then it simulates the lr-oracle perfectly. The probability that A outputs 1 in the real world then, is the probability that b landed on 0 and B output 0 in addition to the probability that b landed on 1 and B output 1. This gives us:

$$\begin{aligned}\mathbf{Adv}^{\text{ind}}_{\Pi}(A,k) &= \Pr[A^{\mathcal{E}(pk,\cdot)}(pk) \Rightarrow 1] - \Pr[A^{\mathcal{E}(pk,0^{|\cdot|})}(pk) \Rightarrow 1] \\ &= \Pr[A^{\mathcal{E}(pk,\cdot)}(pk) \Rightarrow 1] - \frac{1}{2} \\ &= \frac{1}{2} \cdot \Pr[B^{\mathcal{E}(pk,\text{Left}(\cdot,\cdot))}(pk) \Rightarrow 1] + \frac{1}{2} \cdot \Pr[B^{\mathcal{E}(pk,\text{Right}(\cdot,\cdot))}(pk) \Rightarrow 0] - \frac{1}{2} \\ &= \frac{1}{2} \cdot (\Pr[B^{\mathcal{E}(pk,\text{Left}(\cdot,\cdot))}(pk) \Rightarrow 1] - \Pr[B^{\mathcal{E}(pk,\text{Right}(\cdot,\cdot))}(pk) \Rightarrow 1]) \\ &= \frac{1}{2} \cdot \mathbf{Adv}^{\text{lr}}_{\Pi}(A,k)\end{aligned}$$

Problem 22. *Let $\Pi = (\mathcal{K}, \mathcal{E}, \mathcal{D})$ be a public-key encryption scheme. Can it be ind-secure with each of the following "defects"? Briefly justify each answer you give.*

Part A. Encryption of a plaintext M leaks the last bit of M — it is easily computable from the ciphertext C.

Not secure. Consider the following attack. Ask the oracle with $(0^n, 1^n)$. From the ciphertext, if the last bit of the plaintext is 0 then answer 0, otherwise answer 1. It can be easy to see that this attack always outputs the right bit. Hence our advantage is 1.

Part B. Encryption of a plaintext M leaks the length of M — it is easily computable from the ciphertext C.

Could be IND-secure. In fact, we usually imagine that $|M|$ *does* leak, although this certainly isn't "required" in any way. Consider when the scheme Π is derived from an ind-secure public-key encryption by appending the length of the plaintext to the ciphertext. In the ind-experiment, when an adversary makes a query about an n-bit string, regardless of the behavior of the oracle, the adversary will receive the encryption of an n-bit string. Appending the length of the plaintext to the ciphertext therefore does no harm, as the adversary already knows this information before it receives the answer from the oracle.

There is a notion of length-hiding encryption where one is not allowed to leak $|M|$. But length-hiding encryption is impossible on an infinite domain.

Part C. Encryption of a plaintext M leaks the identity of the key pk with which it is encrypted—it is easy to distinguish if a given ciphertext was meant for Alice (it's encrypted under her key) or for Bob (it's encrypted under his).

Could be IND-secure. Consider the case Π is derived from an ind-secure public-key encryption by appending the public key to the ciphertext. In the ind-experiment, when an adversary makes a query, regardless of the behavior of the oracle, the adversary will receive the encryption under the same public key. Appending the public key to the ciphertext therefore does no harm, as the adversary already knows this information before it receives the answer from the oracle.

There is a notion of *anonymity* for public-key encryption that forbids this kind of leakage.

Part D. Encryption of equal-length plaintexts M and M' can take radically different amounts of time.

Could be IND-secure. The standard model does not capture timing, and an oracle is assumed to execute its code in unit time. In practice, this scheme is vulnerable to *timing attack*, which are important but are outside the IND definition.

Part E. Encryption of the secret key sk under its public key pk leaks sk — it is easily computable from the ciphertext C.

Could be IND-secure. Since the adversary lacks the ability to get a session key encrypted, our security notion is effectively silent on what *would* happen if the secret key were public-key encrypted. Here is a more formal argument (more formal than I expected you to get). Consider a scheme $\widetilde{\Pi}$ that is derived from an ind-secure public-key encryption scheme Π, except that encrypting sk will always result in sk. Assume in addition that Π has a "tidiness" property: for all possible (pk, sk) and message M, if C is a valid ciphertext of M then the decryption

of C under any key that is not sk will not output M. More formally, $\forall M\, \forall sk'\, \forall (pk, sk) \twoheadleftarrow \mathcal{K} : ((sk' \neq sk) \Rightarrow \Pr[C \twoheadleftarrow \mathcal{E}_{pk}(M); M' \leftarrow \mathcal{D}_{sk'}(C); M' = M] = 0)$. We remark that this assumption is reasonable if Π is, for example, the hybrid Diffie-Hellman based scheme covered in the lecture. Now suppose that there is an adversary A^f that breaks $\widetilde{\Pi}$. We can construct an adversary B^g that breaks Π as follows. The adversary B computes $c = \mathcal{E}_{pk}(0^n)$ and then runs A as a black box. For each query (x, y) of A, the adversary asks for $g(x, y)$. If $\mathcal{D}_x(c) = 0^n$ then by our tidiness property, x must be sk. Similar statement holds for y. Suppose in this way B finds sk (either x or y), B will output 0 if $\mathcal{D}_{sk}(g(x, y)) = x$, and outputs 1 if $\mathcal{D}_{sk}(g(x, y)) = y$. Otherwise, if neither x nor y is sk, the adversary B will return $g(x, y)$ to A and continue. At the end of the ind-experiment, B will output whatever A outputs. During the experiment, unless A can query $x \equiv sk$ or $y \equiv sk$, the simulation by B will give A the true interaction with Π and thus B wins with advantage $\mathbf{Adv}^{\mathrm{ind}}_{\widetilde{\Pi}}(A)$. However, if A manages to query sk then B's advantage will be $1 \geq \mathbf{Adv}^{\mathrm{ind}}_{\widetilde{\Pi}}(A)$. Hence $\mathbf{Adv}^{\mathrm{ind}}_{\Pi}(B) \geq \mathbf{Adv}^{\mathrm{ind}}_{\widetilde{\Pi}}(A)$.