CMSC 858K — Advanced Topics in Cryptography February 17, 2004

Lecture 7

Lecturer: Jonathan Katz

Scribe(s): *Nagaraj Anthapadmanabhan*
Minkyoung Cho
Ji Sun Shin
Nick L. Petroni, Jr.

1 Introduction

In the last set of lectures, we introduced definitions of adaptively-secure non-interactive zero knowledge and semantically-secure encryption. Based on these, we presented a construction of a public-key encryption scheme secure against *non-adaptive* chosen-ciphertext attacks (CCA1). This encryption scheme was proposed by Naor and Yung in 1990 [3].

In this lecture, we complete the proof of non-adaptive chosen-ciphertext security for the Naor-Yung construction from the previous lecture. Next, we show that the scheme is *not* secure against adaptive chosen-ciphertext attacks by showing a counterexample; we also examine where the proof breaks down. Then, we introduce the definition of a digital signature scheme and the notion of security for a one-time strong signature scheme. Finally, we present a public key encryption scheme secure against *adaptive* chosen-ciphertext attacks (CCA2). This encryption scheme was constructed by Dolev, Dwork, and Naor in 1991 [1].

2 Naor-Yung Construction

The Naor-Yung construction relies on an underlying semantically-secure public-key encryption scheme $(\mathsf{Gen}, \mathcal{E}, \mathcal{D})$ and an adaptively-secure non-interactive zero-knowledge proof system $(\mathcal{P}, \mathcal{V})$. Given these, the scheme is defined as follows:

$\mathsf{Gen}^*(1^k)$: $(pk_1, sk_1) \leftarrow \mathsf{Gen}(1^k)$
$(pk_2, sk_2) \leftarrow \mathsf{Gen}(1^k)$
$r \leftarrow \{0,1\}^{\mathrm{poly(k)}}$
$pk^* = (pk_1, pk_2, r)$
$sk^* = sk_1$

$\mathcal{E}^*_{(pk_1,pk_2,r)}(m)$: pick $w_1, w_2 \leftarrow \{0,1\}^*$
$c_1 \leftarrow \mathcal{E}_{pk_1}(m; w_1)$
$c_2 \leftarrow \mathcal{E}_{pk_2}(m; w_2)$
$\Pi \leftarrow \mathcal{P}\left(r, (c_1, c_2), (m, w_1, w_2)\right)$
output (c_1, c_2, Π)

$\mathcal{D}^*_{sk_1}(c_1, c_2, \Pi)$: if $\mathcal{V}\left(r, (c_1, c_2), \Pi\right) = 0$ then (Verify proof)
    output $\perp$
else
    output $\mathcal{D}_{sk_1}(c_1)$

2.1 CCA1-Security

Theorem 1 *Assuming* $(\mathsf{Gen}, \mathcal{E}, \mathcal{D})$ *is a semantically-secure encryption scheme and* $(\mathcal{P}, \mathcal{V})$ *is an adaptively-secure NIZK proof system, then* $(\mathsf{Gen}^*, \mathcal{E}^*, \mathcal{D}^*)$ *is secure against non-adaptive chosen-ciphertext attacks.*

Recall that in a CCA1 attack an adversary is given access to a decryption oracle before choosing two messages. He does not, however, have access to this oracle after being presented the ciphertext (which he then has to use to guess which message was encrypted). The formal definition is as follows:

Definition 1 An encryption scheme $(\mathsf{Gen}, \mathcal{E}, \mathcal{D})$ is secure against chosen-ciphertext attacks ("CCA1-Secure") if the following is negligible for all PPT algorithms A:

$$\Big|\Pr[(pk, sk) \leftarrow \mathsf{Gen}(1^k);\ (m_0, m_1) \leftarrow A^{\mathcal{D}_{sk}(\cdot)}(pk);\ b \leftarrow \{0,1\};$$

$$C \leftarrow \mathcal{E}_{pk}(m_b);\ b' \leftarrow A(pk, C) :\ b = b'] \Big| - \frac{1}{2}.$$

$\diamondsuit$

We provide here the remainder of the proof of Theorem 1, noting where we left off last lecture. Recall that the idea behind the proof is as follows. We construct a set of games that differ slightly from each other. We show at each step, with the construction of each new game, that the ability of the adversary to distinguish between the two games is negligible (or, in other words, the games are computationally indistinguishable). Transitivity then implies that the first and last games are indistinguishable. But the first game will correspond to the view of an adversary when m_0 is encrypted, while the final game will correspond to the view of an adversary when m_1 is encrypted; thus, this completes the proof.

For the six different games, informal descriptions are as follows:

- Game 0: This is the real game, with the adversary getting an encryption of m_0.
- Game 1: Like Game 0 except that instead of $(\mathcal{P}, \mathcal{V})$, its simulator is used to provide the proof Π.
- Game 2: Like Game 1 except that c_2 is computed as an encryption of m_1.
- Game 2′: Like Game 2 except use sk_2 to decrypt instead of sk_1.
- Game 3: Like Game 2′ except that c_1 is computed as an encryption of m_1.
- Game 3′ Like Game 3 except use sk_1 to decrypt instead of sk_2.
- Game 4: Like Game 3′ except use $(\mathcal{P}, \mathcal{V})$ to provide the proof Π.

Note that Game 4 corresponds exactly to the real game when the adversary gets an encryption of m_1.

Proof of Theorem 1: We first considered the following two games. In each game the adversary is given a pk based on two valid runs of Gen along with a value r. The adversary, with the help of a decryption oracle, then outputs two messages and has to guess which

of those messages was encrypted. The only difference between the two games is that in Game 1 the values for r and Π come from a simulator (and are not a true random string and a real proof, respectively). However, we claim that the probability of the adversary's guess being $b' = 0$ is the essentially same in each case (i.e., only negligibly different).

Game 0 :	Game 1 :
$(pk_1, sk_1), (pk_2, sk_2) \leftarrow \mathsf{Gen}(1^k)$	$(pk_1, sk_1), (pk_2, sk_2) \leftarrow \mathsf{Gen}(1^k)$
$\boxed{r \leftarrow \{0,1\}^{\mathrm{poly}(k)}}$	$\boxed{r \leftarrow \mathsf{Sim}_1(1^k)}$
$pk^* = (pk_1, pk_2, r); \quad sk^* = sk_1$	$pk^* = (pk_1, pk_2, r); \quad sk^* = sk_1$
$(m_0, m_1) \leftarrow A^{\mathcal{D}_{sk^*}(\cdot)}(pk)$	$(m_0, m_1) \leftarrow A^{\mathcal{D}_{sk^*}(\cdot)}(pk)$
$w_1, w_2 \leftarrow \{0,1\}^{poly(k)}$	$w_1, w_2 \leftarrow \{0,1\}^{poly(k)}$
$c_1 \leftarrow \mathcal{E}_{pk_1}(m_0; w_1); \quad c_2 \leftarrow \mathcal{E}_{pk_2}(m_0; w_2)$	$c_1 \leftarrow \mathcal{E}_{pk_1}(m_0; w_1); \quad c_2 \leftarrow \mathcal{E}_{pk_2}(m_0; w_2)$
$\boxed{\Pi \leftarrow \mathcal{P}(r, (c_1, c_2), (m, w_1, w_2))}$	$\boxed{\Pi \leftarrow \mathsf{Sim}_2(c_1, c_2)}$
$b \leftarrow A(c_1, c_2, \Pi)$	$b \leftarrow A(c_1, c_2, \Pi)$

Claim 2 *Let* $\Pr_0[\cdot]$ *and* $\Pr_1[\cdot]$ *represent the probability of event* $\cdot$ *in games 0 and 1 respectively. Then* $|\Pr_0[b = 0] - \Pr_1[b = 0]|$ *is negligible.*

Proof *Summary from last time.* To show that the adversary's choice is not affected by the use of a simulator, the only change from Game 0 to Game 1, we showed that if such a difference *could* be detected by the adversary then the adversary could be used by another adversary, A', to distinguish real proofs from simulated proofs. Using the adversary who can distinguish between games 0 and 1, A' can easily simulate the decryption oracle, obtain messages to encrypt and pass on for a proof, and give the proof to A along with valid ciphertext. The advantage of A' in distinguishing real/simulated proofs is then $|\Pr_0[b = 0] - \Pr_1[b = 0]|$, which is negligible by the security of $(\mathcal{P}, \mathcal{V})$ as an adaptively-secure NIZK proof system. ■

We then made a change to Game 1, where we encrypt message m_1 to get c_2, and called this Game 2.

Game 2 :
$(pk_1, sk_1), (pk_2, sk_2) \leftarrow \mathsf{Gen}(1^k)$
$r \leftarrow \mathsf{Sim}_1(1^k)$
$pk^* = (pk_1, pk_2, r); \quad sk^* = sk_1$
$(m_0, m_1) \leftarrow A^{\mathcal{D}_{sk^*}(\cdot)}(pk)$
$w_1, w_2 \leftarrow \{0,1\}^{poly(k)}$
$c_1 \leftarrow \mathcal{E}_{pk_1}(m_0; w_1); \quad \boxed{c_2 \leftarrow \mathcal{E}_{pk_2}(m_1; w_2)}$
$\Pi \leftarrow \mathsf{Sim}_2(c_1, c_2)$
$b \leftarrow A(c_1, c_2, \Pi)$

Claim 3 *Let* $\Pr_1[\cdot]$ *and* $\Pr_2[\cdot]$ *represent the probability of event* $\cdot$ *in games 1 and 2 respectively. Then* $|\Pr_1[b = 0] - \Pr_2[b = 0]|$ *is negligible.*

Proof *Summary from last time.* As with the previous claim, we wish to show that the existence of an adversary who can distinguish between the two games is an impossibility. This time, instead of the zero-knowledge property, we attacked the semantic security of the

underlying encryption scheme. Specifically, we constructed an adversary A' who wishes to guess which of two messages he generated was encrypted with a given key. To achieve his goal, A' uses the adversary who can distinguish between games 1 and 2 to pick two messages. A' passes these on directly and then gets back an encrypted message that is either m_0 or m_1. Giving this to A as c_2 (note that A' has no problem running the simulator), we see that the advantage of A' is exactly $|\Pr_1[b=0] - \Pr_2[b=0]|$. And this must be negligible by semantic security of the underlying encryption scheme. ■

We next defined an event Fake as the event that A submits (c_1, c_2, Π) to the decryption oracle with $\mathcal{D}_{sk_1}(c_1) \neq \mathcal{D}_{sk_2}(c_2)$, but $\mathcal{V}(r, (c_1, c_2), \Pi) = 1$. This represents the event where the adversary is able to trick the verifier into returning true on a bad pair of ciphertexts (i.e., a pair not in the language). We continued with the claim that the probability of Fake occurring in Game 2 is negligible.

Claim 4 $\Pr_2[\mathsf{Fake}]$ *is negligible.*

Proof Since Fake has to do with the use of the decryption oracle, we note that the only important event from A's perspective at the point he uses the oracle is the generation of pk^*. Furthermore, we note that pk^* is created *identically* in games 1 and 2. Therefore $\Pr_1[\mathsf{Fake}] = \Pr_2[\mathsf{Fake}]$. Last time, we went on to show that $|\Pr_1[\mathsf{Fake}] - \Pr_0[\mathsf{Fake}]|$ is negligible by the zero-knowledge property of the proof system and that $Pr_0[\mathsf{Fake}]$ is negligible because of the soundness of the proof system. Therefore $\Pr_2[\mathsf{Fake}]$ is negligible. ■

At this point, we continue our proof from last time (and begin the new material from this lecture). The proof continues by constructing another game, Game $2'$, which is the same as Game 2 *except* we use sk_2 to decrypt instead of sk_1 (in the obvious way).

Game 2':
$(pk_1, sk_1), (pk_2, sk_2) \leftarrow \mathsf{Gen}(1^k)$
$r \leftarrow \mathsf{Sim}_1(1^k)$
$pk^* = (pk_1, pk_2, r)$; $\boxed{sk^* = sk_2}$
$(m_0, m_1) \leftarrow \boxed{A^{\mathcal{D}_{sk^*}(\cdot)}(pk)}$
$w_1, w_2 \leftarrow \{0,1\}^{poly(k)}$
$c_1 \leftarrow \mathcal{E}_{pk_1}(m_0; w_1)$; $c_2 \leftarrow \mathcal{E}_{pk_2}(m_1; w_2)$
$\Pi \leftarrow \mathsf{Sim}_2(c_1, c_2)$
$b \leftarrow A(c_1, c_2, \Pi)$

Corollary 5 $|\Pr_{2'}[b=0] - \Pr_2[b=0]|$ *is negligible.*

Proof From an adversary's point of view, a difference between Game 2 and Game $2'$ occurs only if the event Fake occurs. This is because, as long as both c_1 and c_2 are encryptions of the same message, decryption using sk_1 or sk_2 will make no difference. It is not hard to see that $\Pr_2[\mathsf{Fake}] = \Pr_{2'}[\mathsf{Fake}]$. Since in either game the event Fake occurs with only negligible probability, the corollary follows. ■

We now modify Game $2'$, encrypting m_1 for both c_1 and c_2, and call this Game 3.

$\underline{\text{Game } 3}$:
$(pk_1, sk_1), (pk_2, sk_2) \leftarrow \mathsf{Gen}(1^k)$
$r \leftarrow \mathsf{Sim}_1(1^k)$
$pk^* = (pk_1, pk_2, r); \quad sk^* = sk_2$
$(m_0, m_1) \leftarrow A^{\mathcal{D}_{sk^*}(\cdot)}(pk)$
$w_1, w_2 \leftarrow \{0,1\}^{poly(k)}$
$\boxed{c_1 \leftarrow \mathcal{E}_{pk_1}(m_1; w_1)}; \quad c_2 \leftarrow \mathcal{E}_{pk_2}(m_1; w_2)$
$\Pi \leftarrow \mathsf{Sim}_2(c_1, c_2)$
$b \leftarrow A(c_1, c_2, \Pi)$

Claim 6 $|\Pr_3[b = 0] - \Pr_{2'}[b = 0]|$ *is negligible.*

Proof We note that the argument used in this proof is similar to the one used in between games 1 and 2. Assume that there exists a PPT adversary A such that $|\Pr_3[b = 0] - \Pr_{2'}[b = 0]|$ is NOT negligible. We can then construct an adversary $A'(pk_1)$ which breaks the semantic security of the underlying encryption scheme, thus generating a contradiction, as follows:

$\underline{A'(pk_1)}$:
$(pk_2, sk_2) \leftarrow \mathsf{Gen}(1^k)$
$r \leftarrow \mathsf{Sim}_1(1^k)$
$pk^* = (pk_1, pk_2, r); \quad sk^* = sk_2$
Run $A^{\mathcal{D}_{sk^*}(\cdot)}(pk)$ until it outputs (m_0, m_1)
Query $\mathcal{E}_{pk_1,b}(m_0, m_1)$ to get c_1
$c_2 \leftarrow \mathcal{E}_{pk_2}(m_1)$
$\Pi \leftarrow \mathsf{Sim}_2(c_1, c_2)$
Output $A(c_1, c_2, \Pi)$

Note that $A'(pk_1)$ is PPT because A is PPT. Also, it is possible for A' to simulate the decryption oracle for A: A' can decrypt using sk_2 since (pk_2, sk_2) is generated locally by A'. Next, we can see that if $c_1 = \mathcal{E}_{pk_1}(m_0)$, then this becomes equivalent to Game 2'; if $c_1 = \mathcal{E}_{pk_1}(m_1)$, then this becomes equivalent to Game 3. So, A' distinguishes encryptions of m_0 from encryptions of m_1 with probability $|\Pr_3[b = 0] - \Pr_{2'}[b = 0]|$, which must be negligible by semantic security of the underlying encryption scheme. ■

We next imagine a game Game 3′ in which we revert back to using sk_1 to decrypt rather than sk_2. As in the proof of indistinguishability between Game 2 and Game 2′, it is not hard to see that Game 3′ is indistinguishable from Game 3.

We next construct another game, Game 4, in which we switch back to *real* proofs from *simulated* proofs.

Game 4 :
$(pk_1, sk_1), (pk_2, sk_2) \leftarrow \mathsf{Gen}(1^k)$
$\boxed{r \leftarrow \{0,1\}^{\mathrm{poly}(k)}}$
$pk^* = (pk_1, pk_2, r); \quad sk^* = sk_1$
$(m_0, m_1) \leftarrow A^{\mathcal{D}_{sk^*}(\cdot)}(pk)$
$w_1, w_2 \leftarrow \{0,1\}^{poly(k)}$
$c_1 \leftarrow \mathcal{E}_{pk_1}(m_1; w_1); \quad c_2 \leftarrow \mathcal{E}_{pk_2}(m_1; w_2)$
$\boxed{\Pi \leftarrow \mathcal{P}(r, (c_1, c_2), (w_1, w_2))}$
$b \leftarrow A(c_1, c_2, \Pi)$

So, Game 4 is basically the actual situation where the adversary gets a *real* encryption of m_1 along with a *real* proof.

Claim 7 $|\mathrm{Pr}_4[b = 0] - \mathrm{Pr}_3[b = 0]|$ *is negligible*

Proof The proof is exactly analogous to that used in studying the transition from Game 0 to Game 1. As there, if an adversary could distinguish between the two games, it could be used by another adversary to distinguish real from simulated proofs. However, this violates the (adaptive) zero-knowledge property of the underlying proof system. ■

From the sequence of preceding claims, we can conclude that $|\mathrm{Pr}_4[b = 0] - \mathrm{Pr}_0[b = 0]|$ is negligible. But since the final game is just the real game when the adversary gets an encryption of m_1, and the original game is just the real game when the adversary gets an encryption of m_0, we see that we have proved that $(\mathsf{Gen}^*, \mathcal{E}^*, \mathcal{D}^*)$ is secure against *non-adaptive* chosen-ciphertext attacks. ■

2.2 CCA2-Security

In this section, we will examine why the Naor-Yung construction is not secure against *adaptive* chosen-ciphertext attacks by giving a counter-example. Recall the formal definition of such attacks:

Definition 2 An encryption scheme $(\mathsf{Gen}, \mathcal{E}, \mathcal{D})$ is secure against adaptive chosen-ciphertext attacks ("CCA2-Secure") if the following is negligible for all PPT algorithms A:

$$\Big|\mathrm{Pr}[(pk, sk) \leftarrow \mathsf{Gen}(1^k);\ (m_0, m_1) \leftarrow A^{\mathcal{D}_{sk}(\cdot)}(pk);\ b \leftarrow \{0,1\};$$

$$C \leftarrow \mathcal{E}_{pk}(m_b);\ b' \leftarrow A^{\mathcal{D}_{sk}(\cdot)}(pk, C) :\ b = b'] \Big| - \frac{1}{2},$$

where A cannot query $\mathcal{D}_{sk}(C)$. ♢

Theorem 8 *The Naor-Yung scheme* $(\mathsf{Gen}^*, \mathcal{E}^*, \mathcal{D}^*)$ *is not secure against adaptive chosen-ciphertext attacks (in general). More precisely, for any semantically-secure encryption scheme* $(\mathsf{Gen}, \mathcal{E}, \mathcal{D})$ *there exists an adaptively-secure NIZK proof system* $(\mathcal{P}, \mathcal{V})$ *such that the resulting Naor-Yung construction is demonstrably insecure against adaptive chosen-ciphertext attacks.*

Proof Let $(\mathcal{P}', \mathcal{V}')$ be any adaptively-secure NIZK proof system. Define the proof system $(\mathcal{P}, \mathcal{V})$ as follows:

$$\underline{\mathcal{P}(r, (c_1, c_2), (w_1, w_2))}:$$
$$\text{Output } \mathcal{P}'(r, (c_1, c_2), (w_1, w_2))|0$$

$$\underline{\mathcal{V}(r, (c_1, c_2), \Pi|b)}:$$
$$\text{Output } \mathcal{V}'(r, (c_1, c_2), \Pi)$$

I.e., we introduce a spurious bit in $\mathcal{P}$ and have $\mathcal{V}$ ignore it. (here, "|" denotes concatenation). It is not hard to show that $(\mathcal{P}, \mathcal{V})$ is also an adaptively-secure NIZK proof system. However, if this new proof system is used in the Naor-Yung construction we can construct an adversary A (making a CCA2 attack) which breaks the encryption as follows:

$\underline{A(pk)}$:
Output (m_0, m_1)
Get back $(c_1, c_2, \Pi|0)$
Submit $(c_1, c_2, \Pi|1)$ to the decryption oracle
Get back m_b

The adversary just modifies the last bit of the challenged ciphertext and submits it to the decryption oracle (note that this is allowed under the definition of CCA2 security). By this method, the adversary will get back the actual message as the last bit was just a spurious bit. So, this construction is not secure against *adaptive* chosen-ciphertext attacks. ■

If we examine the proof of security for the Naor-Yung construction and try to analyze where it breaks down in this case, we see that the $\Pr_2[\mathsf{Fake}]$ is no longer negligible. This is because, if the adversary gets a fake proof (say, $(c_1, c_2, \Pi|0)$), he can construct *another* fake proof by changing just the last bit (i.e., $(c_1, c_2, \Pi|1)$).

We mention in passing that one fix this problem by constructing a proof system satisfying a stronger notion of security: namely, that even when an adversary is given a fake proof, it should be unable to construct a *different* fake proof. See [4, 2] for work along this line. Here, however, we discuss a different method which was historically first.

3 Signature Schemes

For the construction that follows, we will need to notion of a digital signature scheme. Of course, such schemes are also very useful in their own right, and maybe we will return to them later in the course.

Definition 3 A *signature scheme* (over some message space $\mathcal{M}$) is a triple of PPT algorithms $(\mathsf{SigGen}, \mathsf{Sign}, \mathsf{Verify})$ such that:

1. SigGen is a randomized algorithm which outputs a verification key vk and a secret key sk (denoted by $(vk, sk) \leftarrow \mathsf{SigGen}(1^k)$).

2. Sign is a (possibly) randomized algorithm which takes a secret key sk and a message $m \in \mathcal{M}$ and outputs a signature σ (denoted by $\sigma \leftarrow \mathsf{Sign}_{sk}(m)$).

3. Verify is a deterministic algorithm which takes a verification key vk, a message $m \in \mathcal{M}$, and a signature σ and outputs 1 or 0 (denoted by $\mathsf{Verify}_{vk}(m, \sigma)$). A 1 indicates that the signature is *valid* and a 0 indicates that the signature is *invalid*.

We require that for all k, for all (vk, sk) output by $\mathsf{SigGen}(1^k)$, and for all $m \in \mathcal{M}$ we have $\mathsf{Verify}_{vk}(m, \mathsf{Sign}_{sk}(m)) = 1$. $\diamondsuit$

The above merely defines the semantics of signing, but does not give any notion of security. Many such definitions are possible, but we will only require a fairly weak definition of security for the present application (note that this definition of security is too weak for signature schemes used to sign, say, documents).

Definition 4 A signature scheme $(\mathsf{SigGen}, \mathsf{Sign}, \mathsf{Verify})$ is a *one-time, strong signature scheme* if the following is negligible for all PPT adversaries A:

$$\Pr\left[\begin{array}{l}(vk, sk) \leftarrow \mathsf{SigGen}(1^k); m \leftarrow A\,(vk)\,; \sigma \leftarrow \mathsf{Sign}_{sk}(m); \\ (m', \sigma') \leftarrow A\,(vk, \sigma) : \mathsf{Verify}_{vk}\,(m', \sigma') = 1 \wedge (m', \sigma') \neq (m, \sigma)\end{array}\right].$$

$\diamondsuit$

What this means is that, given a signature of a message he chooses, an adversary cannot forge a signature for a different message without knowledge of the secret key. (Also, the adversary cannot even forge a *different* signature on the same message.) While we do not prove it here, it is known that one-time, strong signature schemes exist assuming the existence of one-way functions.

Theorem 9 *If one-way functions exist then one-time, strong signature schemes exist.*

4 Dolev-Dwork-Naor Construction

Danny Dolev, Cynthia Dwork, and Moni Naor [1] constructed an encryption scheme secure against adaptive chosen-ciphertext attacks beginning from any underlying semantically-secure scheme, a one-time, strong signature scheme, and an adaptively-secure NIZK proof system. Their construction is discussed in this section.

Let $(\mathsf{Gen}, \mathcal{E}, \mathcal{D})$ be a *semantically secure* encryption scheme. We construct a new encryption scheme $(\mathsf{Gen}', \mathcal{E}', \mathcal{D}')$ as follows:

$\mathsf{Gen}'(1^k)$: $r \leftarrow \{0,1\}^{\mathrm{poly}(k)}$
for $i = 1$ to k
  for $b = 0$ to 1
    $(pk_{i,b}, sk_{i,b}) \leftarrow \mathsf{Gen}(1^k)$ (Generate $2k$ pairs of keys)

$$pk = \begin{pmatrix} pk_{1,0} & pk_{2,0} & \dots & pk_{k,0} \\ pk_{1,1} & pk_{2,1} & \dots & pk_{k,1} \end{pmatrix}, r$$

$$sk = \begin{pmatrix} sk_{1,0} & sk_{2,0} & \dots & sk_{k,0} \\ sk_{1,1} & sk_{2,1} & \dots & sk_{k,1} \end{pmatrix}$$

${\mathcal{E}'}_{pk}(m)$: $(vk, sk) \leftarrow \mathsf{SigGen}(1^k)$
Let $vk = v_1|v_2|\cdots|v_k$ be the binary representation of vk
  (we assume for simplicity that $|vk| = k$)
for $i = 1$ to k
  $w_i \leftarrow \{0,1\}^*$; $c_i \leftarrow \mathcal{E}_{pk_{i,v_i}}(m; w_i)$
$\Pi \leftarrow \mathcal{P}(r, \vec{C}, (m, \vec{w}))$
  (this is a proof that all ciphertexts correspond to same message)
$\sigma \leftarrow \mathsf{Sign}_{sk}(\vec{C}|\Pi)$
output $vk, \vec{C}, \Pi, \sigma$

${\mathcal{D}'}_{sk}(vk, \vec{C}, \Pi, \sigma)$: if $\mathsf{Verify}_{vk}(\vec{C}|\Pi, \sigma) = 0$ (Verify signature)
  output $\perp$
else
  if $\mathcal{V}(r, \vec{C}, \Pi) = 0$ (Verify proof)
    output $\perp$
  else
    output $\mathcal{D}_{sk_{1,v_1}}(c_1)$

Note that the attack we showed on the Naor-Yung scheme fails here since the attack would require an adversary to forge a signature with respect to vk (which is infeasible). Of course, we need a formal proof to show that the scheme resists *all* adaptive chosen-ciphertext attacks.

Theorem 10 *Assuming* $(\mathsf{Gen}, \mathcal{E}, \mathcal{D})$ *is a* semantically secure *encryption scheme,* $(\mathcal{P}, \mathcal{V})$ *is an adaptively-secure NIZK proof system, and a one-time, strong signature scheme is used, then* $(\mathsf{Gen}', \mathcal{E}', \mathcal{D}')$ *is secure against* adaptive *chosen-ciphertext attacks.*

Proof The proof uses the same structure as that of the Naor-Yung construction. We have a PPT adversary A making an adaptive chosen-ciphertext attack on the encryption scheme. We show that the probability that the adversary will succeed is negligible by constructing a series of games and showing that they are all computationally indistinguishable. We begin by defining our original game, which corresponds to the real encryption scheme when the adversary gets an encryption of m_0:

Game 0 :
$\{(pk_{i,b}, sk_{i,b})\}_{1\leq i\leq k; b\in\{0,1\}} \leftarrow \mathsf{Gen}(1^k)$
$r \leftarrow \{0,1\}^{\mathrm{poly}(k)}$
$pk = \{pk_{i,b}\}_{1\leq i\leq k; b\in\{0,1\}}, r; \quad sk = \{sk_{i,b}\}_{1\leq i\leq k; b\in\{0,1\}}$
$(m_0, m_1) \leftarrow A^{\mathcal{D}'_{sk}(\cdot)}(pk)$
$(vk, sk) \leftarrow \mathsf{SigGen}(1^k)$
for $i = 1$ to k
  $c_i \leftarrow \mathcal{E}_{pk_{i,v_i}}(m_0)$
$\boxed{\Pi \leftarrow \mathcal{P}(r, \vec{C}, (m_0, \vec{w}))}$
$\sigma \leftarrow \mathsf{Sign}_{sk}(\vec{C}|\Pi)$
$b \leftarrow A^{\mathcal{D}'_{sk}(\cdot)}(vk, \vec{C}, \Pi, \sigma)$

We begin by stating a technical lemma. Let Forge be the event that A submits a ciphertext $(vk', \vec{C}', \Pi', \sigma')$ to the decryption oracle with:

- $vk' = vk$
- $(\vec{C}', \Pi', \sigma') \neq (\vec{C}, \Pi, \sigma)$
- $\mathsf{Verify}_{vk}(\vec{C}'|\Pi', \sigma') = 1$

Claim 11 $\Pr_0[\mathsf{Forge}]$ *is negligible.*

This follows from the 1-time strong security of the signature scheme. Details omitted.

The proof of security for the Dolev-Dwork-Naor scheme will be completed in the following lecture. ■

References

[1] D. Dolev, C. Dwork, and M. Naor. *Non-Malleable Cryptography*, proceedings of the twenty-third annual ACM Symposium on Theory of Computing (1991).

[2] Y. Lindell. *A Simpler Construction of CCA2-Secure Public-Key Encryption under General Assumptions*, Eurocrypt 2003: 241-254 (2003).

[3] M. Naor and M. Yung. *Public-key Cryptosystems Provably Secure against Chosen-Ciphertext Attacks*, proceedings of the twenty-second annual ACM Symposium on Theory of Computing (1990).

[4] A. Sahai. *Non-Malleable Non-Interactive Zero Knowledge and Adaptive Chosen-Ciphertext Security*, FOCS 1999: 543-553 (1999).