0368.4162 Foundations of Cryptography January 23, 2020

Lecture 12: Adaptive NIZKs and Chosen Ciphertext Attacks

Lecturer: Nir Bitansky

1 Previously on Foundations of Crypto

We defined and constructed NIZKs. Today, we'll see how to use NIZKs to construct encryption schemes with strong security guarantees. Toward that, we'll first consider a more general notion of adaptive NIZKs for multiple statements, and see how to obtain it.

Let's first recall how we defined NIZKs.

Definition 1.1 (NIZK)**.** *A pair of PPT algorithms (P,V) is a NIZK proof in the CRS model for an NP relation R if they satisfy the following properties:*

1. *Completeness: there exists a polynomial ℓ denoting the length of the common random string such that for every $(x,w)\in R$ we have that:*

$$\Pr_{P,crs\leftarrow\{0,1\}^{\ell(|x|)}}[V(x,crs,\pi)=1:\pi\leftarrow P(x,w,crs)]=1 \ .$$

2. *Adaptive Soundness:*

$$\Pr_{crs\leftarrow\{0,1\}^{\ell(n)}}[\exists\pi,x\in\{0,1\}^n\setminus L,:V(x,crs,\pi)=1]<2^{-n} \ .$$

3. *(Non-Adaptive) Zero Knowledge: there exists a PPT algorithm S such that:*

$$\left\{(crs,P(x,w,crs)):crs\leftarrow U_{\ell(|x|)}\right\}_{(x,w)\in R}\approx_c\{S(x)\}_{(x,w)\in R}$$

In most applications of NIZKs, however, we'll be interested in a stronger ZK definition, which holds for multiple statements x, and even if they're chosen adaptively, after the adversary sees the CRS.

Definition 1.2 (Adaptive NIZK)**.**

3. *Adaptive Zero Knowledge: there exist PPT algorithms (S_1,S_2) such that no n.u. PPT D distinguishes the two following games with advantage greater than negligible:*

   Real:

   (a) $crs\leftarrow\{0,1\}^{\ell(n)}$.

   (b) $D^{P(\cdot,\cdot,crs)}(crs)$, *where D submits queries $(x,w)\in R$.*

   Ideal:

   (a) $(crs,td)\leftarrow S_1(1^n)$.

   (b) $D^{S_2(\cdot,td)}(crs)$, *where D submits queries $(x,w)\in R$, and S_2 only receives x.*

Before we see how to construct such adaptive NIZKs from plain NIZKs, we'll see how to use them to obtain a strong form of encryption.

2 CCA Encryption

We've seen so far how to construct public-key encryption schemes with enhanced functionality — homomorphism. Now we'll see how to obtain public-key encryption schemes with enhanced security known as *security against chosen ciphertext attacks* (CCA). Specially, we're interested in scenarios where the attacker may be able to obtain some *restricted access to the decryption algorithm.* Here the attacker may ask to decrypt ciphertexts of her choice (without knowing what's inside), and the questions is whether security can still be guaranteed for other ciphertexts (ones that haven't been decrypted).

This is not a made up model. In reality, the attacker may sometimes implicitly obtain such access by viewing how a receiver'a responds to certain ciphertexts; indeed, communication very often triggers some action. (To see concrete examples from history, you can read the Wikipedia value of CCA encryption.)

Definition 2.1 (CCA2 security)**.** *A public-key encryption scheme (G, E, D) is said to be CCA secure if no n.u. PPT adversary A wins the following game except with probability $1/2 + \mathrm{negl}(n)$:*

1. *$(pk, sk) \leftarrow G(1^n)$,*
2. *A obtains pk and makes queries to $D_{sk}(\cdot)$ and outputs $m_0, m_1 \in \{0,1\}^n$,*
3. *A obtains $ct^* \leftarrow E_{pk}(m_b)$, for a random $b \leftarrow \{0,1\}$.*
4. *A makes additional queries to $D_{sk}(\cdot)$.*
5. *A outputs a guess b^*.*
6. *A wins if $b = b^*$ and it never queried ct^*.*

The scheme is said to be (only) CCA1 secure if item 4 (post-challenge queries) is eliminated.

Remarks:

1. CCA1 intuitively asks that if the attacker gets access to decryption for a limited time, then this doesn't compromise security of future encryptions, whereas CCA2 asks that security is maintained, even if the attacker has unlimited access to decryption.
2. We defined the above only for public-key schemes. In the secret key setting, achieving CCA security is trivial (make sure you understand why).

2.1 Constructing CCA Encryption: The Naor-Yung Paradigm

It is not hard to see that not any CPA secure scheme is also CCA secure. In particular, homomorphic schemes cannot be CCA secure. In fact, it can be shown that CCA2 security implies non-malleability — namely that it is impossible to change a given ciphertexts into a new different ciphertext for a related message.

We will now describe a general paradigm for transforming CPA encryption into a CCA one using NIZKs.

The construction. Let (G, E, D) be a CPA encryption scheme, and let (P, V) be a NIZK. We construct a CCA scheme as follows:

- $G'(1^n)$:
  - sample two pairs of keys $(sk_0, pk_0), (sk_1, pk_1) \leftarrow G(1^n)$,
  - a CRS $crs \leftarrow \{0,1\}^\ell$, and
  - outputs $(sk', pk') = ((sk_0, crs), (pk_0, pk_1, crs))$.
- $E_{pk'}(m)$:

- sample $ct_b \leftarrow E_{pk_b}(m)$, for both $b \in \{0,1\}$,
- compute a NIZK π, for the fact that ct_0, ct_1 encrypt the same plaintext,
- output $ct' = (ct_0, ct_1, \pi)$.

- $D_{sk'}(ct')$:
  - parse $ct' = (ct_0, ct_1, \pi)$,
  - verify the NIZK π, and if invalid output $\perp$,
  - output $D_{sk_0}(ct_0)$.

Roughly speaking, the intuition behind the scheme is the following: as described the scheme doesn't make any use of sk_1 — the decryption oracle only uses sk_0. In particular, encryptions under pk_1 should remain secure. However, note that as long as the attacker only submits queries (ct_0, ct_1) that encrypts the same thing, then access to $D_{sk_0}(\cdot)$ looks exactly like access to $D_{sk_1}(\cdot)$ — the attacker cannot tell which key is used. The soundness of the NIZK guarantees that this is indeed the case. Accordingly, there is an indistinguishable world where encryptions under pk_0 should remain secure — so actually both should remain secure even in the presence of the decryption oracle. As we shall see, however, formalizing this intuition is a bit tricky.

We'll start by showing that the scheme is CCA1 secure.

Claim 2.2. *The system is CCA1 secure.*

Proof. We will show this by a sequence of hybrids games. Starting from the real game and ending in a games where the attacker's view is independent of the challenge bit b.

Game 0: Real Game

1. $pk' = (sk_0, (pk_0, pk_1, crs)) \leftarrow G'(1^n)$,
2. A obtains pk', makes queries to $D_{sk_0}(\cdot)$, and outputs $m_0, m_1 \in \{0,1\}^n$,
3. A obtains $ct' \leftarrow E_{pk_0}(m_b), E_{pk_1}(m_b), \pi$, for a NIZK π and a random $b \leftarrow \{0,1\}$.
4. A outputs a guess b^*.
5. A wins if $b = b^*$ and it never queried ct'.

Can we change $E_{pk_1}(m_b)$ to an encryption of 0? Not quite yet — note that the NIZK prover uses the randomness for this encryption as a witness. To invoke the security of pk_1, we first need to invoke ZK.

Game 1: Simulated CRS and Challenge Proof

1. $pk' = (sk_0, (pk_0, pk_1, \widetilde{crs})) \leftarrow G'(1^n)$, where $\widetilde{crs}, td \leftarrow S_1(1^n)$
2. A obtains pk', makes queries to $D_{sk_0}(\cdot)$, and outputs $m_0, m_1 \in \{0,1\}^n$,
3. A obtains $ct' \leftarrow E_{pk_0}(m_b), E_{pk_1}(m_b), \widetilde{\pi}$, for a NIZK $\widetilde{\pi} \leftarrow S_2(ct_0, ct_1, td)$ and a random $b \leftarrow \{0,1\}$.
4. A outputs a guess b^*.
5. A wins if $b = b^*$ and it never queried ct'.

Now, we can invoke the security of pk_1; indeed, the attacker's view is independent of sk_1 and the randomness used in $E_{pk_1}(m_b)$.

Game 2: Fake ct_1.

1. $pk' = (sk_0, (pk_0, pk_1, \widetilde{crs})) \leftarrow G'(1^n)$, where $\widetilde{crs}, td \leftarrow S_1(1^n)$
2. A obtains pk', makes queries to $D_{sk_0}(\cdot)$, and outputs $m_0, m_1 \in \{0,1\}^n$,
3. A obtains $ct' \leftarrow E_{pk_0}(m_b), E_{pk_1}(0), \widetilde{\pi}$, for a NIZK $\widetilde{\pi} \leftarrow S_2(ct_0, ct_1, td)$ and a random $b \leftarrow \{0,1\}$.
4. A outputs a guess b^*.
5. A wins if $b = b^*$ and it never queried ct'.

Next, we would like to apply the same argument for ct_2. For this, we first move to decrypting using sk_0 and argue that the attacker's view is indistinguishable.

Game 3: Decrypting with sk_1.

1. $pk' = (sk_0, (pk_0, pk_1, \widetilde{crs})) \leftarrow G'(1^n)$, where $\widetilde{crs}, td \leftarrow S_1(1^n)$
2. A obtains pk', makes queries to $D_{sk_1}(\cdot)$, and outputs $m_0, m_1 \in \{0,1\}^n$,
3. A obtains $ct' \leftarrow E_{pk_0}(m_b), E_{pk_1}(0), \widetilde{\pi}$, for a NIZK $\widetilde{\pi} \leftarrow S_2(ct_0, ct_1, td)$ and a random $b \leftarrow \{0,1\}$.
4. A outputs a guess b^*.
5. A wins if $b = b^*$ and it never queried ct'.

We can claim that the last two hybrids are indistinguishable provided that the attacker doesn't make decryption queries (ct_0^*, ct_1^*, π^*) such that the ciphertext do not encrypt the same message and yet π^* is accepted by the NIZK verifier. This follows from (1) soundness, (2) ZK. Indeed, had the crs been real, rather than simulated, then there wouldn't be such false proof. Moreover, since the simulated CRS, $\widetilde{crs}$, is indistinguishable from a real one, false proofs are hard to find also relative to the simulated CRS.

Now, we can invoke the security of pk_0.

Game 4: Fake ct_0.

1. $pk' = (sk_0, (pk_0, pk_1, \widetilde{crs})) \leftarrow G'(1^n)$, where $\widetilde{crs}, td \leftarrow S_1(1^n)$
2. A obtains pk', makes queries to $D_{sk_1}(\cdot)$, and outputs $m_0, m_1 \in \{0,1\}^n$,
3. A obtains $ct' \leftarrow E_{pk_0}(0), E_{pk_1}(0), \widetilde{\pi}$, for a NIZK $\widetilde{\pi} \leftarrow S_2(ct_0, ct_1, td)$ and a random $b \leftarrow \{0,1\}$.
4. A outputs a guess b^*.
5. A wins if $b = b^*$ and it never queried ct'.

Finally, note that in game 4, the attacker's view is independent of b, and thus wins with probability $1/2$. □

Is the scheme CCA2? Imagine that we add post-challenge decryption queries. Does the above hybrid strategy still hold? This is actually not clear. The problem is that once the attacker sees a simulated proof $\widetilde{\pi}$, it may be that soundness no longer holds. In particular, imagine that the NIZK had a bit, which the verifier anyhow ignores. Then, the attacker could take the proof given in game 2 with the challenge cipher, flip this bit, and get a decryption query allowing it to tell whether decryption is done using sk_0 or sk_1.

We want an extra property, which says that a simulated proof (for a possibly false statement), cannot be mauled into a new accepting proof for a different false statement. This is called *simulation soundness* [].

Definition 2.3 (Simulation Soundness)**.** *An adaptive NIZK is called one-time simulation sound if for any n.u. PPT prover P^* wins the following game with negligible probability:*

1. $(crs, td) \leftarrow S_1(1^n)$,
2. $P^*(crs)$ *outputs* x *and obtains* $\pi \leftarrow S_2(x, td)$.
3. P^* *wins if it outputs* $x' \notin L \cup x$ *along with an accepting proof.*

It is not hard to see that one-time simulation soundness is enough to extend the proof we've seen to show that the scheme is CCA2 secure.

Claim 2.4 ([?])**.** *Any NIZK can be turned into one that is simulation sound.*

3 How to Construct Adaptive NIZKs

We will now discuss how to turn any NIZK as in Definition 1.1 into one that is adaptive ZK.

4 A Detour: Witness Indistinguishability and ZAPs

We've seen that completely non-interactive ZK is impossible without a trusted setup. What about weaker notions of privacy? One such notion commonly considered is that of witness indistinguishability, which says that a proof shouldn't reveal which one of possibly many witnesses was used to produce the proof.

We will show that any NIZK (with non-adaptive ZK) implies witness indistinguishable proof in the random string model, where only the verifier has to trust the crs. These are called ZAPs [?].

Definition 4.1 (ZAPs)**.** *A proof system* (P, V) *is a ZAP for an NP relation* R *if it satisfies:*

1. *Completeness: there exists a polynomial* ℓ *denoting the length of the verifier random string such that for every* $(x, w) \in R$ *we have that:*

$$\Pr_{P, r \leftarrow \{0,1\}^{\ell(|x|)}}[V(x, r, \pi) = 1 : \pi \leftarrow P(x, w, r)] = 1 \ .$$

2. *Adaptive Soundness:*

$$\Pr_{r \leftarrow \{0,1\}^{\ell(n)}}[\exists \pi, x \in \{0,1\}^n \setminus L, : V(x, r, \pi) = 1] < 2^{-n} \ .$$

3. *Witness Indistinguishability:*

$$\{P(x, w_0, r)\}_{r, (x,w_0), (x,w_1) \in R} \approx_c \{P(x, w_1, r)\}_{r, (x,w_0), (x,w_1) \in R} \ .$$

Note that the prover does not need to trust the random string in the sense that WI is guaranteed for any choice of string r. Also, note that WI against adaptive statements is guaranteed — the above holds for any statement, in particular, one that depends arbitrarily on r.

An adaptive NIZK from ZAPs and OWFs. Given ZAPs and OWFs gives a simple construction of adaptive NIZKs in the CRS model. In what follows let $G : \{0,1\}^n \to \{0,1\}^{2n}$ be PRG.

- The CRS includes (r_1, r_2) where $r_1 \in \{0,1\}^\ell$ is a random verifier string for a ZAP, and $r_2 \in \{0,1\}^{2n}$.
- To prover a statement $(x, w) \in R$, the prover gives a ZAP (relative to r_1) attesting that:

$$\text{"either } x \in L(R) \text{ or } r_2 = G(s) \text{ for some } s \in \{0,1\}^n\text{,"}$$

  where the prover uses w as the witness.

The idea behind soundness is that if r_2 is chosen truly at random, then except with probability 2^{-n}, r_2 is not in the image of G, and thus the (adaptive) soundness of the ZAP kicks in. The zero knowledge simulator $S = (S_1, S_2)$ works as follows. $S_1(1^n)$ generates $r_2 = G(s)$, and saves s as a trapdoor (r_1 can be chosen at random). $S_2(x, s)$ proves the ZAP using s as the witness. Adaptive zero knowledge follows from the (adaptive) witness indistinguishability of the ZAP.

ZAPs from NIZK (with Non-Adaptive ZK): the Dwork-Naor Construction. Let (P, V) be a NIZK with a CRS of length ℓ, we construct a ZAP as follows:

- The verifier random string, includes ℓ strings $r_1, \ldots, r_\ell$ each $\leftarrow \{0,1\}^\ell$.
- The prover picks $r \leftarrow \{0,1\}^\ell$. It proves the statement with respect to each one of the strings $r \oplus r_i$.

Soundness is shown by a union bound. For any fixed r the probability that there exists a false statement relative to all of the strings $r \oplus r_i$ is at most $2^{-n\ell}$. Thus by a union bound, the system is sound. Witness indistinguishability is shown by a hybrid argument. Here we use the fact that the NIZK is a (non-adaptive) WI (note that the statement depends on verifier string $r_1, \ldots, r_\ell$, but not on the prover string r.

The construction is similar in spirit to Lautmann's proof that $BPP \subset \Sigma_2$, and is sometimes called *reverse randomization.*

References