CR09 – Homework 2022-2023
Due date Jan. 25, 2023, 11:59pm

1 NIZK for Honest Partial Decryption

Recall that in the e-voting scheme constructed in class, the parties jointly compute a ciphertext (c_0, c_1), which is an encryption of the sum of their individual votes, with respect to the public key $h = \prod_{i=1}^{N} h_i$, where $h_i = g^{s_i}$ (s_i is party P_i's locally generated share of the secret key s). All parties broadcast the partial decryption $d_i = c_0^{s_i}$, and the result is computed as $\mathsf{DLOG}_g(c_1 / \prod_{i=1}^{N} d_i)$.

Question 1. Construct a NIZK for each party P_i which guarantees that the decrypted result is indeed the plaintext of (c_0, c_1) with respect to the public key h.

Question 2. Given all encrypted votes $(c_0^{(i)}, c_1^{(i)})_{i \leq N}$, how could P_i prove *anonymously* to an external player that they took part to the vote (*i.e.* without revealing which party P_i they are)? A high-level description of how to construct the appropriate zero-knowledge proof will suffice here.

2 The Hybrid Proof Technique

The purpose of this exercise it to introduce the hybrid proof technique, a standard method to prove reductions between indistinguishability-based security notions in cryptography.

We recall that a (bit) *commitment scheme* is an interactive protocol between two parties, a *sender* S and a *receiver* R, with common input 1^λ. The sender has an input bit b. The commitment is

- (computationally) *hiding* if for every PPT receiver R^*, it holds that
$$\{\langle S(0), R^* \rangle (1^\lambda)\}_{\lambda \in \mathbb{N}} \overset{\mathsf{comp}}{\equiv} \{\langle S(1), R^* \rangle (1^\lambda)\}_{\lambda \in \mathbb{N}},$$
where $\overset{\mathsf{comp}}{\equiv}$ denotes computational indistinguishability;
- (perfectly) *binding* if the probability, over the coin tosses of R, that a transcript T can simultaneously satisfy $T = \langle S(0; r_0), R^* \rangle (1^\lambda)$ and $T = \langle S(1; r_1), R^* \rangle (1^\lambda)$ (where r_0, r_1 denote some choice of random tape for S) is negligible. That is, if S reveal that they committed to a bit b (by revealing their random tape), they cannot pretend to have committed $1 - b$ instead.

We say further that the scheme is *n-vector hiding* (where $n = n(\lambda)$ is some integer polynomial in λ) if for every PPT R^* and every pairs of n-tuples $\bar{b} = (b_1, \cdots, b_n) \in \{0,1\}^n$ and $\bar{b}' = (b'_1, \cdots, b'_n) \in \{0,1\}^n$, it holds that
$$\{\langle \bar{S}(\bar{b}), \bar{R}^* \rangle (1^\lambda)\}_{\lambda \in \mathbb{N}} \overset{\mathsf{comp}}{\equiv} \{\langle \bar{S}(\bar{b}'), \bar{R}^* \rangle (1^\lambda)\}_{\lambda \in \mathbb{N}},$$
where $\langle \bar{S}(\bar{b}), \bar{R}^* \rangle$ (resp. $\bar{b}'$) denote the interactive protocol obtained by running the n instances $\langle S(b_i), R^* \rangle$ (resp. b'_i) of the commitment, for $i = 1$ to n, in parallel.

Question 1. Fix $i \leq n$. Let $\bar{b}$ and $\bar{b}'$ bit two vectors that differ only in their i-th coordinate. Prove formally that if the scheme is *hiding*, then it must hold that
$$\{\langle \bar{S}(\bar{b}), \bar{R}^* \rangle (1^\lambda)\}_{\lambda \in \mathbb{N}} \overset{\mathsf{comp}}{\equiv} \{\langle \bar{S}(\bar{b}'), \bar{R}^* \rangle (1^\lambda)\}_{\lambda \in \mathbb{N}}. \tag{1}$$

Hint: assume the existence of a PPT distinguisher $\mathcal{A}$ with advantage ε in distinguishing the distribution ensemble of Equation 1. Given $\mathcal{A}$, construct a distinguisher $\mathcal{B}$ with advantage ε against the hiding property of the commitment scheme.

Question 2. Fix arbitrary n-bit vectors $\bar{b}$ and $\bar{b}'$. Describe a sequence of $n+1$ *hybrid scenarios* $H_0, \cdots, H_n$ (also called *games*), where H_0 corresponds to the scenario where S commits to $\bar{b}$, and H_n to the scenario where S commits to $\bar{b}'$, such that for any $i < n$, if there exists a PPT distinguisher $\mathcal{A}$ with advantage ε in distinguishing the distribution ensemble of H_i from H_{i+1}, then there exists a distinguisher $\mathcal{B}$ with advantage ε against the hiding property of the commitment scheme.

Question 3. Conclude by proving formally that the n-vector hiding property is implied by the hiding property. The reduction should construct, given any adversary $\mathcal{A}$ with non-negligible advantage ε against the n-vector hiding property, an adversary $\mathcal{B}$ with non-negligible advantage ε' against the hiding property. The quantity ε/ε' is called the *loss* of the reduction. What is the loss of this reduction?

Note. This reduction finishes the proof, seen in class, that 3-coloring is zero-knowledge under the assumption that the commitment scheme is hiding. Combined with the existence of hiding commitment schemes from any pseudorandom generator, which in turn exist assuming any one-way function, this yields (via extended Karp reduction) a proof that CZK contains NP.

Question 4. We introduce the following *parallel composition lemma* for witness indistinguishability (recall that being witness-indistinguishable is a property of the *prover* for a *language relation* $\mathcal{R}_{\mathcal{L}}$):

Lemma 1. *Let $\mathcal{L}$ be an NP language with relation $\mathcal{R}_{\mathcal{L}}$, P be a PPT algorithm, and assume that $\langle P, V\rangle$ is a witness-indistinguishable interactive proof for $\mathcal{R}_{\mathcal{L}}$. Let $n = n(\cdot)$ be a polynomial, and let P_n denote the program which, on common input $(x_1, \cdots, x_n)$ and private input $(w_1, \cdots, w_n)$, runs P in parallel n times where the i-th copy invoke P on common input x_i and private input w_i. Then P_n is witness-indistinguishable for the parallel relation $\bar{\mathcal{R}}_{\bar{\mathcal{L}}} = \{(\bar{x}, \bar{w}) \ : \ \forall i, (x_i, w_i) \in \mathcal{R}_{\mathcal{L}}\}$.*

Using the same approach as in Questions 1 to 3, prove the parallel composition lemma.

Note. We saw in class that (1) 3COL is zero-knowledge assuming OWF, and (2) zero-knowledge implies witness indistinguishability. Therefore, the parallel repetition of 3COL is witness indistinguishable by the above lemma, which yields a 4-round WI with negligible soundness error for 3COL (two rounds for the PRG-based commitment, three for the ZK proof itself, the first round of the proof being merged with the second round of the commitment), hence for all NP languages, assuming any OWF.

3 More on the Power of Interactive Proofs

The purpose of this exercise will be to prove two additional statements about the power of the classes CZK and IP: (1) IP supports efficient amplification, allowing to reduce the soundness error at an exponential rate, and (2) assuming OWFs, $\mathsf{CZK} = \mathsf{IP} = \mathsf{PSPACE}$ (recall that we have seen in class that $\mathsf{IP} = \mathsf{PSPACE}$ and that, under OWFs, CZK contains NP; here, we will show that the power of CZK goes in fact all the way to IP).

Question 1. Recall that AM is the class of languages that admit a two-round public-coin interactive proof: Arthur (the verifier) sends a random string r to Merlin, and Merlin replies with a single message m, from which Arthur accepts or rejects. Using the fact that all languages in NP admit a (computational) zero-knowledge proof, show that if there exists a computationally hiding and perfectly binding commitment scheme, $\mathsf{AM} \subset \mathsf{CZK}$.

Hint: the verifier's decision to accept or reject is a deterministic polynomial-time function of the transcript (r, m). Use this observation to build, for a language $\mathcal{L} \in \mathsf{AM}$, a well-chosen NP-language $\mathcal{L}'$ that depends both on the commitment and on an AM protocol for $\mathcal{L}$. The zero-knowledge proof for AM should use a ZK proof for this language (which exists since $\mathsf{NP} \subset \mathsf{CZK}$).

Question 2. Building upon the ideas of question 1, and the interactive protocol constructed in class for quantified boolean formulae, prove that assuming a commitment scheme, QBF is in CZK. Since QBF is PSPACE-complete, this implies $\mathsf{CZK} = \mathsf{IP} = \mathsf{PSPACE}$.

Hint: you don't need to use any specific property of the protocol, beyond the fact that it is an interactive proof, and that is is public coin. *Use these properties to build a well-chosen NP-language, and conclude as in question 1.*

Question 3. Denote of $\mathsf{IP}_{\alpha,\beta}$ the class of languages admitting an interactive proof system with completeness α and soundness β. That is,

$$x \in \mathcal{L} \Rightarrow \exists P / \Pr[\mathsf{out}_V \langle V, P \rangle(x) = 1] \geq \alpha \ ,$$

and

$$x \notin \mathcal{L} \Rightarrow \forall P^*, \Pr[\mathsf{out}_V \langle V, P^* \rangle(x) = 1] \leq \beta \ .$$

In particular, we have $\mathsf{IP} = \mathsf{IP}_{2/3,1/3}$. Show that:

1. $\mathsf{IP}_{1-2^{-|x|},1-2^{-|x|}} = \mathsf{IP}$.
2. $\mathsf{IP}_{1,1/3} = \mathsf{IP}$.
3. $\mathsf{IP}_{2/3,1} = \mathsf{NP}$.

4 Security of Schnorr's Signature Scheme

Let $\mathbb{G}$ denote a cyclic group of prime order p and let g denote a generator of $\mathbb{G}$. We remind Schnorr's signature scheme. Let $H : \{0,1\}^* \to \mathbb{Z}_p$ denote a hash function, modeled as a random oracle in the rest of the exercise.

- $\mathsf{Setup}(\mathbb{G}, g)$: Sample $s \leftarrow U(\mathbb{Z}_p)$ and let $\mathsf{sk} \leftarrow s, \mathsf{vk} \leftarrow g^s$. Return $(\mathsf{vk}, \mathsf{sk})$.
- $\mathsf{Sign}(\mathsf{sk}, m \in \{0,1\}^*)$: Sample $r \leftarrow U(\mathbb{Z}_p)$. Let $t \leftarrow g^r$, $c \leftarrow H(\mathsf{vk}, m, t)$, and $z \leftarrow r + c \cdot \mathsf{sk}$. Return $\sigma = (t, c, z)$.
- $\mathsf{Verify}(\mathsf{vk}, m, \sigma)$: Parse σ as (t, c, z). Return 1 if:
  (i) $c = H(\mathsf{vk}, m, t)$;
  (ii) $g^z = t \cdot \mathsf{vk}^c$
  Else return 0.

The goal of this exercise is to prove that Schnorr's signature is secure in the random oracle model assuming the discrete logarithm problem is hard in $\mathbb{G}$.

We also recall that a signature is secure if any PPT adversary $\mathcal{A}$ fails to winning the following game, except with negligible probability.

The challenger first sample $(\mathsf{vk}, \mathsf{sk}) \leftarrow \mathsf{Setup}(\mathbb{G}, g)$ and sends vk to the adversary $\mathcal{A}$. The adversary can adaptively make sign queries to the challenger by sending a message $m_i \in \{0,1\}^*$. In response, the challenger returns σ_i where $\sigma_i \leftarrow \mathsf{Sign}(\mathsf{sk}, m_i)$. Let q_S denote the number of sign queries made by the adversary, and denote by $L = ((m_1, \sigma_1), \ldots, (m_{q_S}, \sigma_{q_S}))$ the list of queries/answers. The adversary can further make hash queries to the challenger to learn the evaluation of H on an input of its choice. Let q_H denote the number of hash queries made by the adversary, and assume without loss of generality that all such queries are of the form (vk, m_i, t_i), for $m_i \in \{0,1\}^*$ and $t_i \in \mathbb{G}$.

The adversary wins if it returns a pair $(m^*, \sigma^*) \notin L$ such that $\mathsf{Verify}(\mathsf{vk}, m^*, \sigma^*) = 1$.

Question 1. Let $\sigma^* = (t^*, c^*, z^*)$ denote the output of $\mathcal{A}$. Show that if $\mathcal{A}$ has not made a hash query for (vk, m^*, t^*), then $\mathcal{A}$ wins with probability at most $1/p$.

Question 2. Show that, the challenger (which controls the random oracle) can answer signing queries correctly without using the secret key. That is, show that there exists a poly-time simulator Sim, which given vk and a query m_i, produces a valid signature $\sigma_i = (t_i, c_i, z_i)$.

Question 3. Show that, given (t^*, c^*, z^*) and (t^*, c'^*, z'^*) such that each satisfy Equation (ii) of the verification process and $c'^* \neq c^*$, one can compute sk.

Question 4. Conclude about security of Schnorr's signature by relying on the splitting lemma.

5 5-Round Zero-Knowledge Proofs

We consider Pedersen commitment scheme, defined over a cyclic group $\mathbb{G}$ of prime order p with 2 random public generators (g, h) as follows: to commit to a message $m \in \mathbb{Z}_p$, the sender samples a randomness $r \leftarrow U(\mathbb{Z}_p)$ and produces a commitment $g^m h^r \leftarrow \mathsf{Com}(m; r)$. To open a commitment, it reveals (m, r).

Question 1. Show that Pedersen commitment is perfectly hiding and computationally binding. That is, for $m \in \mathbb{Z}_p$, let $\mathcal{D}_m$ the distribution $\{\mathsf{Com}(m; r) \mid r \leftarrow U(\mathbb{Z}_p)\}$, then show that $\Delta(\mathcal{D}_m, \mathcal{D}_{m'}) = 0$, for all m, m'. Also, show that a sender being able to open a commitment C to distinct messages $m' \neq m$ can be used to solve the discrete logarithm problem in $\mathbb{G}$.

Question 2. Let Σ denote a sigma-protocol with challenge space $\mathbb{Z}_p$ and achieving special-soundness and special honest-verifier zero-knowledge. Show that such a protocol achieves soundness $1/p$.

Question 3. Consider the 5-round protocol defined by adding two additional rounds before running the sigma-protocol:

1. P first sends public parameters for Pedersen commitment $(\mathbb{G}, g, h)$ to V;
2. V replies with a group element $C \leftarrow \mathsf{Com}(c; r)$, with $c, r \leftarrow U(\mathbb{Z}_p)$, which is a commitment to a uniformly random challenge c;
3. This is the first round of the sigma-protocol Σ;
4. V reveals the challenge by opening C. P aborts if the opening is invalid;
5. This is the third round of the sigma-protocol Σ.

Show that this protocol now achieves (computational) zero-knowledge (and not just honest-verifier zero-knowledge), and still has soundness $1/p$.