CS 65500
Advanced Cryptography

Lecture 16: Zero-Knowledge Proofs - II

Instructor: Aarushi Goel

Spring 2025

Agenda

- → Zero-Knowledge Proof for Graph Isomorphism
- → Proofs of Knowledge

Defining Interactive Proofs (without Zero-Knowledge)

Definition: A protocol π between a prover P and a verifier V is an interactive proof system for a language L if V is a PPT machine and the following properties hold:

- Completeness: $\forall x \in L$

$$\Pr[\mathrm{Out}_V[P(x) \longleftrightarrow V(x)] = 1] = 1$$

- Soundness: There exists a negligible function $\nu(.)$, s.t., $\forall x \notin L$, $\forall \lambda \in \mathbb{N}$ and all adversarial provers P^*,

$$\Pr[\mathrm{Out}_V[P^*(x) \longleftrightarrow V(x)] = 1] \leq \nu(\lambda)$$

We can also modify the above definition to consider PPT provers

Proofs that are only secure against PPT provers are called arguments.

Formalizing the Notion of Zero-Knowledge.

→ If an interactive proof convinced the verifier that $x \in L$, then this interactive proof should not leak any information about the witness w that the prover used to participate in in the interactive proof.

→ In other words, whatever the verifier saw during the interactive proof, it could have *simulated* on its own using x, L and the fact that $x \in L$.

Defining Zero-Knowledge

Definition: An interactive proof π between P & V for a language L with witness relation R is said to be zero-knowledge if for every n.u. PPT adversary V^*, there exists a (expected) PPT simulator S, such that $\forall x \in L, \forall w \in R(x), \forall\ z \in \{0,1\}^*$ and $\forall \lambda \in \mathbb{N}$, the following two distributions are computationally indistinguishable:

1. $\{ \text{View}_{V^*} [P(x,w) \longleftrightarrow V^*(x,z)] \}$
2. $\{ S^{V^*}(1^\lambda, x, z, L) \}$

We can also consider the notions of statiscal/perfect zero-knowledge against unbounded adversaries, if the above distributions are statistically close (or identical respectively)

Graph Isomorphism

→ Let $G=(V,E)$ be a graph, wher V is the set of vertices & E is the set of edges

→ $G_0=(V_0,E_0)$ and $G_1=(V_1,E_1)$ are said to be isomorphic if there exists a permutation π, s.t.

* $V_1=\{\pi(v) \mid v\in V_0\}$

* $E_1=\{(\pi(v_1),\pi(v_2)) \mid (v_1,v_2)\in E_0\}$

in other words, $G_1=\pi(G_0)$

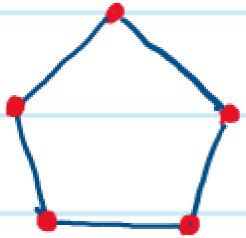

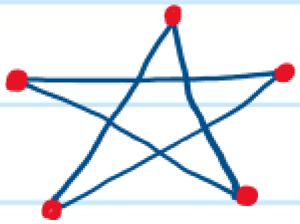

Graph isomorphism is in NP.

Zero-Knowledge Proof for Graph Isomorphism

Prover wants to convince the verifier that graphs G_0 and G_1 are isomorphic without revealing π, where $\pi(G_0) = G_1$

Prover

(G_0, G_1, π)

Verifier

(G_0, G_1)

Sample a random permutation σ

Sample a random bit $b \in \{0,1\}$

$H = \sigma(G_b)$

$\xrightarrow{H}$

Samples a random bit $b' \in \{0,1\}$

$\xleftarrow{b'}$

$$\tau = \begin{cases} \sigma & \text{if } b' = b \\ \sigma \cdot \pi^{-1} & \text{if } b=0,\ b'=1 \\ \sigma \cdot \pi & \text{if } b=1,\ b'=0 \end{cases}$$

$\xrightarrow{\tau}$

Output 1 iff $H = \tau(G_{b'})$

Completeness: If $G_1 = \pi(G_0)$ then it will always be the case that $\tau(G_{b'}) = \sigma(G_b)$

Soundness: If the verifier is honest, it chooses b' randomly. In this case, if $\nexists\, \pi$ s.t. $G_1 = \pi(G_0)$, then $\tau(G_{b'}) = \sigma(G_b)$ iff $b = b'$. This only happens with probability $1/2$.

(We can repeat this protocol multiple times and let the verifier accept only if the protocol outputs 1 on each repitition. If we repeat λ times, the probability that the output will be 1 on each repitition is $\frac{1}{2^{\lambda}}$

Zero-Knowledge: $\text{View}_{V^*}[P(G_0, G_1, \pi) \longleftrightarrow V^*(G_0, G_1)] = (b', H, \sigma)$

$S^{V^*}(1^\lambda, G_0, G_1)$:

→ Randomly choose a permutation σ & $b \xleftarrow{\$} \{0,1\}$.

→ Set $H = \sigma(G_b)$ & $b' = V^*(G_0, G_1, H)$

→ If $b' = b$ output (b', H, σ), otherwise restart with a new σ, b.

In order to show that S is a valid simulator, it suffices to prove that if G_0 and G_1 are isomorphic, then

1. S runs in expected polynomial time

2. Distribution of its output is indistinguishable from View_{V^*}

1. **S runs in expected polynomial time:** Since σ is random, b' cannot depend on b. Therefore b & b' are chosen independently
$\Rightarrow$ $b = b'$ with probability $1/2$
$\Rightarrow$ S must run twice in expectation before haulting.

2. **Distribution of its output is indistinguishable from View_{V^*}:**
Since we have argued that $\Pr[b = b'] = 1/2$, this implies that whether or not S halts on a choice of (b, σ) is independent of (b, σ) and therefore of H.
$\Rightarrow$ (b', H, σ) are distributed identically to View_{V^*}

Zero Knowledge Proofs for NP

→ This was an exampu of a ZKP with perfect zero-knowledge.

→ In general we know of ZKPs with computational zero-knowledge for NP-complete languages such as circuit SAT, Graph-hamiltonicity etc.

⇒ There exist computational ZKPs *for* all languages in NP.

Proofs of Knowledge

→ Soundness in a zero-knowledge proof ensures that if $x \notin L$, then a malicious prover will not be able to compute an accepting proof with high probability.

→ In some applications, however, we require a stronger guarantee.

→ In particular, we want that even if $x \in L$, but if the prover does not know the corresponding witness w, s.t., $R_L(x, w) = 1$, then he cannot compute an accepting proof with high probability

→ In other words, if a prover can compute an accepting proof to prove that $x \in L$, then with a high probability, he must know a corresponding witness, s.t., $R_L(x, w) = 1$

→ This property is called *knowledge soundness* and proofs that satisfy this property are called *proofs of knowledge*.

→ This property is formalized by showing existence of an *extractor* algorithm which given oracle access to the adversarial prover can extract a valid witness corresponding to the statement

Defining Zero-Proofs of Knowledge

Definition: A zero-knowledge proof π between P & V for a language L, with witness relation R_L is said to be a proof of knowledge with knowledge error ϵ, if $\exists$ an algorithm E^{P^*}, called an extractor, that runs in expected polynomial time, such that the following holds for every x and every P^*

$$\Pr[\, \mathrm{Out}_V [P^*(x) \longleftrightarrow V(x)] = 1] - \Pr[\, R_L(x, w) = 1 \quad w \leftarrow E^{P^*}(x)] \leq \epsilon$$

ZKPs that only satisfy knowledge soundness against PPT provers are called arguments of knowledge

Knowledge Soundness in ZKP for Graph Isomorphism.

Consider an extractor EXT that proceeds as follows:

1. EXT queries the malicious prover P^* to get the first round message H
2. EXT then queries P^* on input $b'=0$ to get a third round message τ_0
3. EXT again queries P^* on input $b'=1$ to get another third round message τ_1
4. Given τ_0 and τ_1, EXT can now learn π, such that $\pi(G_0) = G_1$

How?