Foundations of Probabilistic Proofs

A course by **Alessandro Chiesa**

Lecture 05

Zero-Knowledge IPs

Zero Knowledge IPs

paper that first introduced the notion of zero knowledge →

The Knowledge Complexity of Interactive Proof Systems

Shafi Goldwasser	Silvio Micali	Charles Rackoff
MIT	MIT	University of Toronto



Benefits of interaction and randomness so far:

- capture many languages beyond NP (coNP, $P^{\#P}$, PSPACE)
- delegate computation (bounded-depth circuits)

Today we study another benefit: ZERO KNOWLEDGE.

Informally, we seek IPs that protect the privacy of the honest prover.
The honest prover should reveal no information beyond the necessary bit "$x \in L$".

We illustrate this notion via the language $GI = \{(G_0, G_1) \mid G_0 \equiv G_1\}$.
Recall that GI is in NP: the witness is any isomorphism between the graphs.
Hence there is a trivial IP: the IP prover sends an isomorphism to the IP verifier.

CHALLENGE: what if the isomorphism is a private input of the honest prover?

How to design an alternative IP for GI (achieving completeness and soundness)
where the honest prover reveals no information beyond $G_0 \equiv G_1$?

Interactive Proofs for Relations

A relation is a set of instance-witness pairs $R = \{(x,w) : \dots\}$.

The corresponding language is $L(R) := \{x : \exists\, w \text{ s.t. } (x,w) \in R\}$.

Languages can be viewed as relations with empty witnesses: $L = \{x : \dots\}$, $R = \{(x,\perp) : \dots\}$

Example:
- GI as a language $L_{GI} = \{(G_0,G_1) : G_0 \equiv G_1\}$.
- GI as a relation $R_{GI} = \{((G_0,G_1),\sigma) : G_0 = \sigma(G_1)\}$. Note that $L_{GI} = L(R_{GI})$.

The definition of an IP directly extends from languages to relations.

def: (P,V) is an IP for a relation R with completeness error ε_c and soundness error ε_s this holds:

① completeness:

$$\forall\,(x,w)\in R \quad \Pr_{s_p,s_v}\left[\langle P(x,w;s_p), V(x;s_v)\rangle = 1\right] \geq 1-\varepsilon_c$$

② soundness:

$$\forall\, x\notin L(R) \quad \forall\, \tilde{P} \quad \Pr_{s_v}\left[\langle \tilde{P}, V(x;s_v)\rangle = 1\right] \leq \varepsilon_s$$

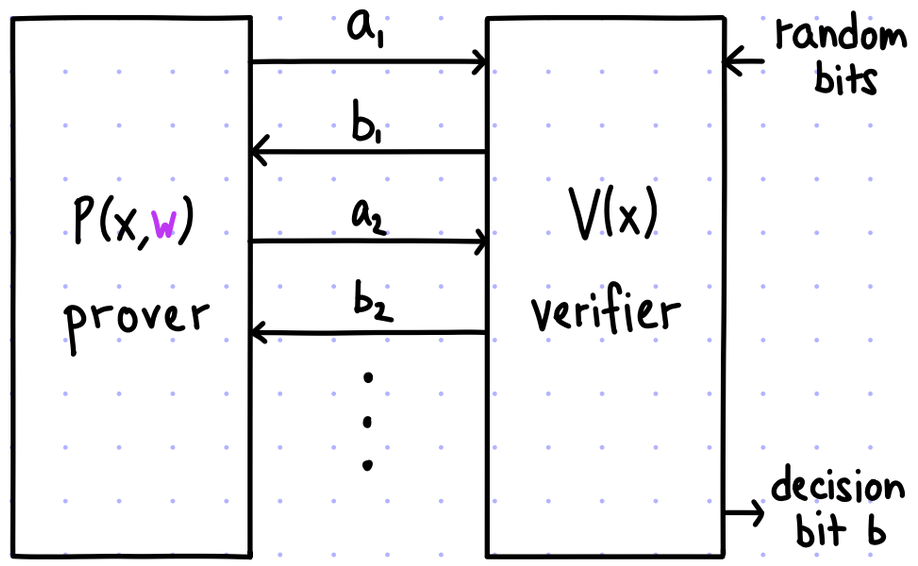

Today we focus on the (more general) case of IPs for relations.

Zero Knowledge against Honest Verifiers

An IP (P,V) for a relation R is **honest-verifier zero knowledge** (HVZK) if

$\exists$ polynomial-time probabilistic algorithm S (known as the **simulator**) such that

$$\forall (x,w)\in R \quad S(x) \equiv \mathrm{View}(P,V,x,w)$$

($\equiv$ means equality as distributions)

Here $\mathrm{View}(P,V,x,w) := (\rho, x, a_1, \ldots, a_k)$ is everything V sees when interacting with $P(x,w)$:

- V's randomness ρ,
- V's input instance x,
- P's messages $a_1, \ldots, a_k$.

INTERPRETATION:

The honest verifier **can simulate the interaction by itself**, without talking to the honest prover.

The simulator captures this by sampling the honest verifier's view.

In particular, $\forall$ function f $\quad \forall (x,w)\in R \quad f(S(x)) \equiv f(\mathrm{View}(P,V,x,w))$.

NOTES:

- HVZK is a **joint property** of the honest prover P & honest verifier V.
  (This is like the completeness property, also a joint property of P and V.)
- HVZK is preserved under **sequential and parallel repetition** of the IP.

Honest-Verifier ZK for Graph Isomorphism [1/2]

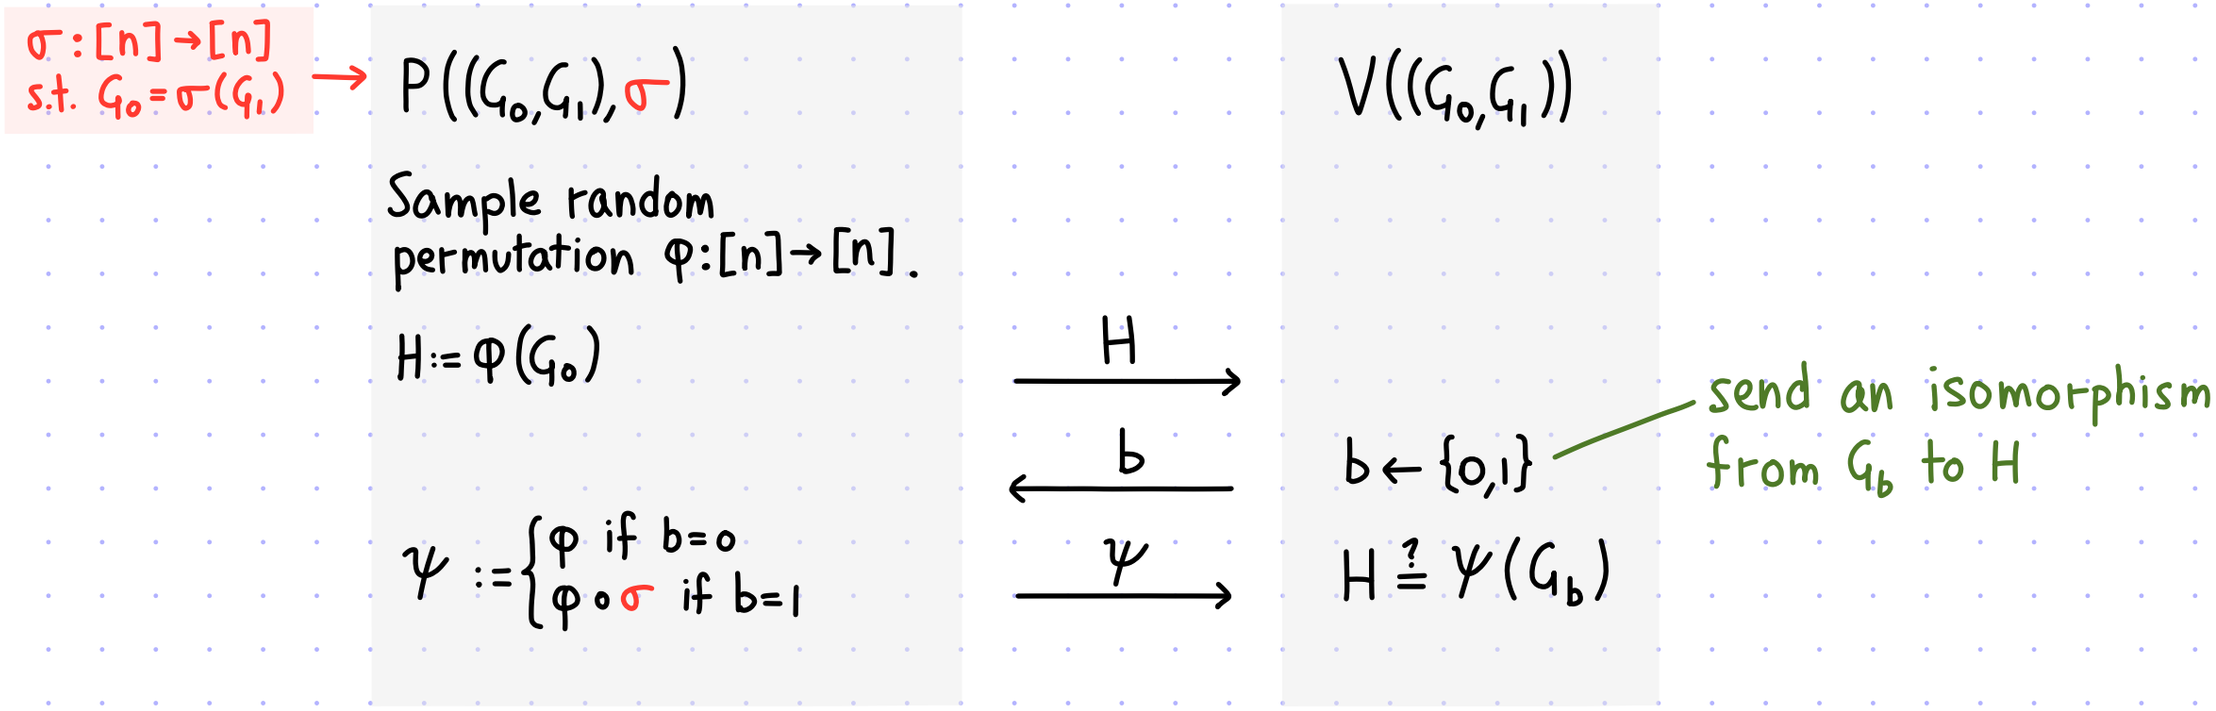

First we argue that this is an IP for GI.

Completeness: Suppose that $((G_0,G_1),\sigma) \in R_{GI}$ (i.e. $\sigma:[n]\to[n]$ is s.t. $G_0 = \sigma(G_1)$).

If $b=0$: the prover sends $\psi := \varphi$, so the test $H \stackrel{?}{=} \psi(G_0)$ passes since $H := \varphi(G_0)$.

If $b=1$: the prover sends $\psi := \varphi\circ\sigma$, so the test $H \stackrel{?}{=} \psi(G_1)$ passes since $(\varphi\circ\sigma)(G_1) = \varphi(G_0)$ and $H := \varphi(G_0)$.

Soundness: Suppose that $(G_0,G_1) \notin L_{GI} = L(R_{GI})$.

Then H (prover's first message) can be isomorphic to at most one of G_0 and G_1.

Any malicious prover gets caught w.p. $\geq 1/2$.

Honest-Verifier ZK for Graph Isomorphism [2/2]

claim: (P,V) is HVZK

proof: Fix $((G_0,G_1),\sigma) \in R_{GI}$.

The honest verifier's view is

$$((G_0,G_1),H,b,\psi)$$

where
- H equals $\psi(G_b)$
- b is a random bit
- ψ is a random permutation on $[n]$ (independent of b)

$P((G_0,G_1),\sigma)$

Sample random permutation $\varphi:[n]\to[n]$.

$H := \varphi(G_0)$

$$\psi := \begin{cases} \varphi & \text{if } b=0 \\ \varphi\circ\sigma & \text{if } b=1 \end{cases}$$

H →

← b

ψ →

$V((G_0,G_1))$

$b \leftarrow \{0,1\}$

$H \stackrel{?}{=} \psi(G_b)$

Consider the following polynomial-time probabilistic algorithm:

$S((G_0,G_1))$:=
1. Sample $b \leftarrow \{0,1\}$.
2. Sample random permutation $\psi:[n]\to[n]$.
3. Compute $H := \psi(G_b)$.
4. Output $((G_0,G_1),H,b,\psi)$.

Since $G_0 \equiv G_1$, the output of S is equidistributed as V's view. ■

Zero Knowledge against Malicious Verifiers [1/3]

We can strengthen zero knowledge to require that even verifiers $\tilde{V}$ that deviate from the prescribed protocol cannot learn any information besides the bit "$x \in L(R)$".

How does the simulator S know about the malicious verifier $\tilde{V}$?

restricting the simulator strengthens the property

- existential simulation: $\forall$ efficient $\tilde{V}$ $\exists$ efficient $S_{\tilde{V}}$ $\forall (x,w) \in R$ $S_{\tilde{V}}(x) \equiv View(P, \tilde{V}, x, w)$
- universal simulation: $\exists$ efficient S $\forall$ efficient $\tilde{V}$ $\forall (x,w) \in R$ $S(\tilde{V}, x) \equiv View(P, \tilde{V}, x, w)$
- black-box simulation: $\exists$ efficient S $\forall$ efficient $\tilde{V}$ $\forall (x,w) \in R$ $S^{\tilde{V}}(x) \equiv View(P, \tilde{V}, x, w)$

Note that malicious-verifier ZK is a property of the honest prover P alone.
(Compare with: completeness is of P&V; soundness is of V; HVZK if of P&V.)

REMARK: Preserving malicious-verifier ZK under repetition of the IP is tricky.

- Sequential repetition preserves auxiliary-input malicious-verifier ZK.
  The condition is strengthened to $\begin{cases} \forall aux\ S_{\tilde{V}}(x, aux) \equiv View(P, \tilde{V}_{(aux)}, x, w) & \text{for existential simulation} \\ \forall aux\ S(\tilde{V}, x, aux) \equiv View(P, \tilde{V}_{(aux)}, x, w) & \text{for universal simulation} \end{cases}$
  Black-box simulation supports auxiliary inputs as is.
- Parallel repetition does NOT, in general, preserve malicious-verifier ZK (assuming plausible crypto).
  This is even for black-box simulation.

Zero Knowledge against Malicious Verifiers [2/3]

We focus on BLACK-BOX simulation:

$$\exists \text{ efficient } S \;\; \forall \text{ efficient } \tilde{V} \;\; \forall (x,w) \in R \quad S^{\tilde{V}}(x) \equiv \text{View}(P, \tilde{V}, x, w)$$

What is "efficient"?

Ideally: $\tilde{V}$ and S are polynomial-time probabilistic algorithms.

Problem: theorem [Barak Lindell 2002]:

If L has an IP with round complexity $k = O(1)$, soundness error $\varepsilon_s = \text{negl}(n)$, and $S^{\tilde{V}}(x)$ runs in polynomial time (for polynomial-time $\tilde{V}$) then $L \in BPP$.

Common workaround (there are others): S runs in EXPECTED polynomial time.

RECALL: Let A be a probabilistic algorithm. For any input x, $\text{time}(A(x))$ is a random variable.

- (STRICT) polynomial time means $\exists c$ s.t. $\text{time}(A(x)) \leq |x|^c$.
- EXPECTED polynomial time means $\exists c$ s.t. $\mathbb{E}[\text{time}(A(x))] \leq |x|^c$.

E.g. $\text{time}(A(x))$ equals $|x|$ w.p. $1 - \frac{1}{2^{|x|}}$ and $2^{|x|}$ w.p. $\frac{1}{2^{|x|}}$ has $\mathbb{E}[\text{time}(A(x))] = |x| - \frac{|x|-1}{2^{|x|}} \leq |x|$.

Zero Knowledge against Malicious Verifiers [3/3]

An IP (P,V) for a relation R is malicious-verifier zero knowledge (MVZK) if

$\exists$ expected polynomial-time probabilistic algorithm S (called the simulator) such that

$\forall$ polynomial-time probabilistic $\tilde{V}$ $\forall (x,w) \in R$ $S^{\tilde{V}}(x) \equiv View(P, \tilde{V}, x, w)$

This definition can be achieved (non-trivially).

We see an example in the next slide.

Even so, one can establish LIMITATIONS ON ROUND COMPLEXITY:

theorem: Suppose L has a k-round IP with $\varepsilon_s = negl(n)$ and (expected polynomial-time) simulation.

- If $k=2$ then $L \in BPP$ (even for existential simulation). ← [Oren 1987][Goldreich Oren 1993]
- If $k=3$ and simulation is black box then $L \in BPP$. ↙ [Goldreich Krawczyk 1990]
- If $k=O(1)$, the IP is public-coin, and simulation is black box then $L \in BPP$.

These motivate the study of NON-BLACK-BOX simulation.

Malicious-Verifier ZK for Graph Isomorphism

claim: (P,V) is MVZK

proof: Fix $((G_0,G_1),\sigma) \in R_{GI}$ and $\tilde{V}$.

The view of the malicious verifier $\tilde{V}$ is

$$((G_0,G_1), H, \tilde{b}, \psi)$$

where
- H equals $\psi(G_{\tilde{b}})$
- $\tilde{b}$ is distributed as $\tilde{V}(H)$
- ψ is a random permutation on $[n]$ (independent of $\tilde{b}$)

$P((G_0,G_1),\sigma)$		$V((G_0,G_1))$
Sample random permutation $\varphi:[n]\to[n]$.		
$H := \varphi(G_0)$	$\xrightarrow{H}$	
	$\xleftarrow{b}$	$b \leftarrow \{0,1\}$
$\psi := \begin{cases} \varphi & \text{if } b=0 \\ \varphi\circ\sigma & \text{if } b=1 \end{cases}$	$\xrightarrow{\psi}$	$H \stackrel{?}{=} \psi(G_b)$

Consider the following EXPECTED polynomial-time probabilistic algorithm:

$S^{\tilde{V}}((G_0,G_1)) :=$
1. Sample $b \leftarrow \{0,1\}$.
2. Sample random $\psi:[n]\to[n]$.
3. Compute $H := \psi(G_b)$.
4. Give H to $\tilde{V}$ to get $\tilde{b}$.
5. If $\tilde{b} \neq b$ then GOTO 1.
6. Output $((G_0,G_1), H, \tilde{b}, \psi)$.

↑ S uses $\tilde{V}$ only as a black-box

$S^{\tilde{V}}((G_0,G_1))$ runs in expected polynomial time:

$G_0 \equiv G_1 \to H$ is independent of b
$\to \tilde{b}$ is independent of b
$\to Pr[\tilde{b}=b] = 1/2 \to \mathbb{E}[\#\text{rewinds}] = 2$.

$S^{\tilde{V}}((G_0,G_1))$ follows the desired distribution:

$$Pr[\tilde{b}=0 \mid \tilde{b}=b] = \frac{Pr[\tilde{b}=0 \wedge \tilde{b}=b]}{Pr[\tilde{b}=b]} = \frac{Pr[\tilde{b}=0]\cdot 1/2}{1/2} = Pr[\tilde{b}=0]$$

■

Limitations of Zero Knowledge

What happens more generally?

def:
- HVZK-IP = all languages that have IPs with honest-verifier zero knowledge
- (MV)ZK-IP = all languages that have IPs with malicious-verifier zero knowledge

Straightforward: $BPP \subseteq$ (simulator does nothing) $MVZK\text{-}IP \subseteq HVZK\text{-}IP \subseteq IP$.

Moreover, we proved that $GI \in MVZK\text{-}IP$ (and GI is not known to be in BPP).

Q: What languages have zero Knowledge IPs?

theorem: $HVZK\text{-}IP \subseteq AM \cap coAM$

Hence we do not expect that $NP \subseteq HVZK\text{-}IP$. (Since $NP \subseteq coAM$ directly implies that $coNP \subseteq IP[k=O(1)]$, which implies that the Polynomial Hierarchy collapses [Boppana, Håstad, Zachos 1987].)

So far we discussed PERFECT ZERO KNOWLEDGE (PZK), where $S(x)$ equals the verifier's view.
The above limitation holds even for honest-verifier STATISTICAL ZERO KNOWLEDGE (SZK),
which relaxes the requirement on the simulator for the honest verifier:
require only that $S(x)$ and $View(P,V,x,w)$ are statistically close.

Intuition on the Limits of HVZK-IP

Suppose that (P,V) is an HVZK IP for R.

Let S be the HVZK simulator. We know that $\forall (x,w)\in R \quad S(x) \equiv View(P,V,x,w)$.

Q: What does $S(x)$ do if $x \notin L(R)$?

① $S(x)$ outputs a view $(\varsigma, x, a_1, \ldots, a_k)$ that is REJECTING (with non-negligible probability)

② $S(x)$ outputs a view $(\varsigma, x, a_1, \ldots, a_k)$ that is ACCEPTING (but for a negligible probability)

If option ① then $L(R)\in BPP$ (in the weaker "infinitely often" sense): use the simulator to decide.

↳ $\exists$ BPP machine that decides L on infinitely many (maybe not all) input sizes

So suppose that option ② holds.

OBSERVATION: $x\in \overline{L(R)} \rightarrow S(x)$ and $View(P,V,x,w)$ are statistically far (for every w).

Indeed, soundness implies that $View(P,V,x,w)$ is accepting with small probability for every w.

There are public-coin $O(1)$-round IPs for showing that $S(x)$ and $View(P,V,x,w)$ are close ($x\in L$) & far ($x\notin L$).

lemma: $L \in HVZK\text{-}IP \rightarrow L \in AM[k=O(1)]$ [Aiello Håstad 1991]

lemma: $L \in HVZK\text{-}IP \rightarrow \bar{L} \in AM[k=O(1)]$ [Fortnow 1987]

Implies that $HVZK\text{-}IP \subseteq AM \cap coAM$

because $AM[k=O(1)] = AM[k=1]$ [Babai 1985].

Example: from HVZK-IP for GI to IP for GNI

Consider the HVZK IP for GI and its simulator:

$P_{GI}((G_0,G_1),\sigma)$

Sample random permutation $\varphi:[n]\to[n]$.

$H := \varphi(G_0)$

$\psi := \begin{cases} \varphi & \text{if } b=0 \\ \varphi\circ\sigma & \text{if } b=1 \end{cases}$

H →

← b

ψ →

$V_{GI}((G_0,G_1))$

$b \leftarrow \{0,1\}$

$H \stackrel{?}{=} \psi(G_b)$

$S_{GI}((G_0,G_1)) :=$

1. Sample $b \leftarrow \{0,1\}$.
2. Sample random permutation $\psi:[n]\to[n]$.
3. Compute $H := \psi(G_b)$.
4. Output $((G_0,G_1), H, b, \psi)$.

We proved that if $G_0 \equiv G_1$ then $S_{GI}((G_0,G_1)) \equiv \text{View}(P_{GI}, V_{GI}, (G_0,G_1), \sigma)$.

If $G_0 \not\equiv G_1$ then $S_{GI}((G_0,G_1))$ still outputs accepting views but with a DIFFERENT distribution.

We can use this to recover the protocol for GNI (the complement of GI)!

$P_{GNI}((G_0,G_1))$

find $\tilde{b}$ s.t. $H \equiv G_{\tilde{b}}$

H ←

$\tilde{b}$ →

$V_{GNI}((G_0,G_1))$

$b \leftarrow \{0,1\}$

$\pi \leftarrow \{\text{permutations on vertices}\}$

$H := \pi(G_b)$

$\tilde{b} \stackrel{?}{=} b$

V_{GNI} runs the simulator for (P_{GI}, V_{GI}) and then asks the prover to show that $G_0 \not\equiv G_1$ by requesting a guess for the bit b. Indeed H determines b when $G_0 \not\equiv G_1$, H is independent from b when $G_0 \equiv G_1$.

IPs with Computational Zero Knowledge

We still want zero knowledge for NP (and more). What to do?

One approach is COMPUTATIONAL ZERO KNOWLEDGE:
relax the requirement on the simulator to

$S^{\tilde{V}}(x)$ and $\text{View}(P,\tilde{V},x,w)$ are computationally close

$\{A_x\}_{x\in S}$ and $\{B_x\}_{x\in S}$ s.t.
$\forall$ poly-size circuit family $\{D_n\}_{n\in\mathbb{N}}$
$|\Pr[D_{|x|}(A_x)=1] - \Pr[D_{|x|}(B_x)=1]| = \text{negl}(|x|)$

This leads to corresponding complexity classes: HVCZK-IP & MVCZK-IP

theorem: if OWFs exist then MVCZK-IP = IP

one-way functions

We sketch a weaker result:

theorem: commitment schemes → NP ⊆ MVCZK-IP

Everything Provable is Provable in Zero-Knowledge

Michael Ben-Or	Hebrew University
Oded Goldreich	Technion – Israel Institute of Technology
Shafi Goldwasser	M.I.T. Laboratory for Computer Science
Johan Håstad	Royal Institute of Technology, Sweden
Joe Kilian	M.I.T. Laboratory for Computer Science
Silvio Micali	M.I.T. Laboratory for Computer Science
Phillip Rogaway	M.I.T. Laboratory for Computer Science

Proofs that Yield Nothing But Their Validity or All Languages in NP Have Zero-Knowledge Proof Systems

ODED GOLDREICH SILVIO MICALI AND AVI WIGDERSON

The limitations of [Goldreich Krawczyk 1990] and [Barak Lindell 2002] hold even for CZK.

Circumventing the limitations motivates the study of non-black-box universal simulators.

The GMW Protocol for 3COL [1/3]

Consider the NP-complete 3COL (graph 3-coloring) problem:

- $L_{3COL} := \{ G=(V,E) : G \text{ is a 3-colorable graph} \}$
- $R_{3COL} := \{ (G,a) : a:V\to[3] \text{ is a 3-coloring of } G=(V,E) \}$
  $\forall (i,j)\in E.\ a_i \neq a_j$

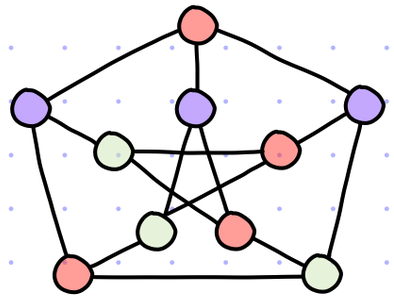

3-coloring of the Petersen graph

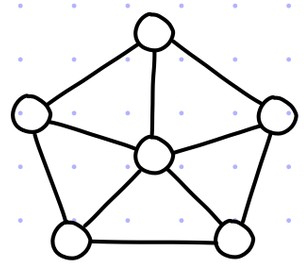

not 3-colorable

We study the Goldreich-Micali-Wigderson (GMW) protocol for graph 3-colorability.

It is an MVCZK-IP for R_{3COL}. This yields MVCZK-IPs for all of NP.

MAIN TOOL: COMMITMENT SCHEMES (for simplicity, non-interactive)

A tuple $CM = (CM.Commit, CM.Check)$ that satisfies these properties:

- completeness: $\forall m\in\mathcal{M}\ \ Pr[CM.Check(cm,m,pf)=1 \mid (cm,pf)\leftarrow CM.Commit(m)]=1$.
- perfect binding: $\forall cm\in\mathcal{C}\ \ |\{ m\in\mathcal{M} : \exists pf\in\mathcal{O} \text{ s.t. } CM.Check(cm,m,pf)=1 \}|=1$.
- computational hiding:

$$\forall m_0,m_1\in\mathcal{M},\ \{ cm_0 \mid (cm_0,pf_0)\leftarrow CM.Commit(m_0)\} \overset{c}{\equiv} \{ cm_1 \mid (cm_1,pf_1)\leftarrow CM.Commit(m_1)\}$$

(computationally close)

EXAMPLE: El Gamal commitment scheme

$CM.Setup(1^\lambda)\to(\mathbb{G},g,h)$ — group of prime order q; random group elements in $\mathbb{G}$

$CM.Commit(m\in\mathbb{G})\to(cm,r)$ where $cm := (g^r, m\cdot h^r)\in\mathbb{G}^2$ and r is random in $\mathbb{Z}_q$

$CM.Check(cm,m,r) := cm \overset{?}{=} (g^r, m\cdot h^r)$

Note:
in every commitment scheme, hiding or binding must be computational

The GMW Protocol for 3COL

We describe the GMW protocol for graph 3-colorability.

Sequential repetition of the GMW protocol reduces soundness error while preserving MVZK.

$P(G, a: V \to [3])$

Sample random permutation $\varphi: [3] \to [3]$

Permute colors: $b := \varphi \circ a$

$\forall v \in V$, $(cm_v, pf_v) \leftarrow CM.Commit(b_v)$.

$\xrightarrow{(cm_v)_{v \in V}}$

$\xleftarrow{(i,j)}$

$\xrightarrow{(b_i, pf_i, b_j, pf_j)}$

$V(G)$

$(i,j) \leftarrow E$

$b_i, b_j \in [3] \quad b_i \neq b_j$

$CM.Check(cm_i, b_i, pf_i) \stackrel{?}{=} 1$
$CM.Check(cm_j, b_j, pf_j) \stackrel{?}{=} 1$

This protocol is an IP for R_{3COL}.

- completeness error $\varepsilon_c = 0$: If a is a 3-coloring of G then, for every permutation φ, b is also a 3-coloring of G. Hence, $\forall (i,j) \in E$, b_i and b_j are distinct colors in $[3]$.
- soundness error $\varepsilon_s = 1 - \frac{1}{|E|}$: Fix a malicious IP prover $\tilde{P}$. Let $(\tilde{cm}_v)_{v \in V}$ be its commitments. By perfect binding of CM, $(\tilde{cm}_v)_{v \in V}$ defines a partial coloring $\tilde{a}: V \to [3]$. Since G is not 3-colorable, $\exists (i^*, j^*) \in E$ s.t. $a_{i^*} = a_{j^*}$ (or one of a_{i^*} or a_{j^*} is undefined). If V sends (i^*, j^*) then $\tilde{P}$ cannot convince V to accept.

The GMW Protocol for 3COL [3/3]

lemma: the GMW protocol for R_{3COL} satisfies CZK.

We describe the simulator and only sketch its analysis.

Fix a 3-colorable graph G and a malicious IP verifier $\tilde{V}$.

$P(G, a: V \to [3])$

Sample random permutation $\varphi: [3] \to [3]$

Permute colors: $b := \varphi \circ a$

$\forall v \in V$, $(cm_v, pf_v) \leftarrow CM.Commit(b_v)$.

$\xrightarrow{(cm_v)_{v \in V}}$

$\xleftarrow{(i,j)}$

$\xrightarrow{(b_i, pf_i, b_j, pf_j)}$

$V(G)$

$(i,j) \leftarrow E$

$b_i, b_j \in [3] \quad b_i \neq b_j$

$CM.Check(cm_i, b_i, pf_i) \stackrel{?}{=} 1$
$CM.Check(cm_j, b_j, pf_j) \stackrel{?}{=} 1$

$S^{\tilde{V}}(G) :=$

1. Sample $(i,j) \leftarrow E$.
2. Sample $b_i, b_j \leftarrow [3]$ s.t. $b_i \neq b_j$.
3. $\forall v \in V \setminus \{i,j\}$, set $b_v := 1$.
4. $\forall v \in V$, $(cm_v, pf_v) \leftarrow CM.Commit(b_v)$.
5. Give $(cm_v)_{v \in V}$ to $\tilde{V}$ to get $(\tilde{i}, \tilde{j})$.
6. If $(\tilde{i}, \tilde{j}) \neq (i,j)$ then GOTO 1.
7. Output $(G, (cm_v)_{v \in V}, (\tilde{i}, \tilde{j}), (b_{\tilde{i}}, pf_{\tilde{i}}, b_{\tilde{j}}, pf_{\tilde{j}}))$.

EASY: the output of $S^{\tilde{V}}(G)$, if it halts, follows the desired distribution.

HARD: Does $S^{\tilde{V}}(G)$ run in expected polynomial-time (or even halt)? Computational hiding of CM implies that $\tilde{V}$ cannot "force" $(\tilde{i}, \tilde{j}) \neq (i,j)$ too often. Arguing this is delicate.

Zero Knowledge Beyond IPs

Zero knowledge can be defined for other models of probabilistic proof.

The capabilities and limitations of zero knowledge are (very) different in each setting.

Example: zero knowledge IA

An interactive argument (IA) is an IP whose soundness is relaxed to computational soundness (consider only malicious provers that are efficient).

theorem: OWFs → NP ⊆ MVZK-IA

Idea: modify the GMR protocol to use CM that is perfectly hiding & computationally binding.

Example: Pedersen commitment scheme

Example: zero knowledge MIP

A multi-prover interactive proof (MIP) is a generalization of an IP where the verifier interacts with multiple non-communicating provers.

theorem: MVZK-MIP = MIP

Multi-Prover Interactive Proofs:
How to Remove Intractability Assumptions

Michael Ben-Or* Hebrew University | Shafi Goldwasser* MIT | Joe Kilian‡ MIT | Avi Wigderson§ Hebrew University

Cryptography is replaced by a physical assumption (the provers cannot communicate).

One ingredient of the theorem: unconditional commitments in the MIP model.

Bibliography

Zero-knowledge

- [GMR 1985]: The knowledge complexity of interactive proof-systems, by Shafi Goldwasser, Silvio Micali, Charles Rackoff.
- [GMW 1991]: Proofs that yield nothing but their validity or all languages in NP have zero-knowledge proof systems, by Oded Goldreich, Silvio Micali, Avi Wigderson.
- [Goldreich 2010]: Zero-knowledge: a tutorial, by Oded Goldreich.
- (▸Computer scientist explains zero knowledge proofs in 5 levels of difficulty), by Amit Sahai.
- (▸History of ZK), by Shafi Goldwasser.
- [Ilango 2025]: Gödel in cryptography: effectively zero-knowledge proofs for NP with no interaction, no setup, and perfect soundness, by Rahul Ilango.

Power of ZK

- [BGGHKMR 1988]: Everything provable is provable in zero-knowledge, by Michael Ben-Or, Oded Goldreich, Shafi Goldwasser, Johan Håstad, Joe Kilian, Silvio Micali, Phillip Rogaway.
- [BGKW 1988]: Multi-prover interactive proofs: how to remove intractability assumptions, by Michael Ben-Or, Shafi Goldwasser, Joe Kilian, Avi Wigderson.

Limitations of ZK

- [Fortnow 1987]: The complexity of perfect zero-knowledge, by Lance Fortnow.
- [AH 1987]: Perfect zero-knowledge languages can be recognized in two rounds, by William Aiello, Johan Håstad.
- [GO 1994]: Definitions and properties of zero-knowledge proof systems, by Oded Goldreich, Yair Oren.
- [GK 1998]: On the composition of zero-knowledge proof systems, by Oded Goldreich, Hugo Krawczyk.
- [BL 2002]: Strict polynomial-time in simulation and extraction, by Boaz Barak, Yehuda Lindell.