Foundations of Probabilistic Proofs

A course by **Alessandro Chiesa**

Lecture 22

Public vs Private Coins & Perfect Completeness

Public Coins vs Private Coins

Randomness in interactive proofs comes in different forms.

Ex 1: in the 2-message IP for GNI, the verifier's random bit b must be secret.

Ex 2: in the poly(n)-message IP for TQBF, all verifier randomness is sent to the prover.

TODAY: How do these settings compare?

def: A verifier V is public-coin if its every message is a freshly sampled uniform random string of a prescribed length. (If V has no restrictions then V is private-coin.)

def: AM[k]/MA[k] are languages decidable via k-round public-coin IPs where the verifier/prover moves first. ("A" stands for Arthur = verifier & "M" stands for Merlin = prover)

Trivial: $\forall k$, $AM[k], MA[k] \subseteq IP[k]$.

Surprising: theorem: $\forall k$, $IP[k] \subseteq AM[k+1]$

We study a special case of the theorem today.

Private Coins versus Public Coins in Interactive Proof Systems

Shafi Goldasswer
MIT

Michael Sipser
MIT

Revisiting Graph Non-Isomorphism

theorem: $GNI \in AM[k=1]$ (Previously we proved that $GNI \in IP[k=1]$.)

IDEA: look at graph isomorphism in a quantitative way

def: The automorphism group of a graph $G=(V,E)$ is

$$aut(G) = \{ \pi : V \to V \mid \pi \text{ is a permutation and } \pi(G) = G \} .$$

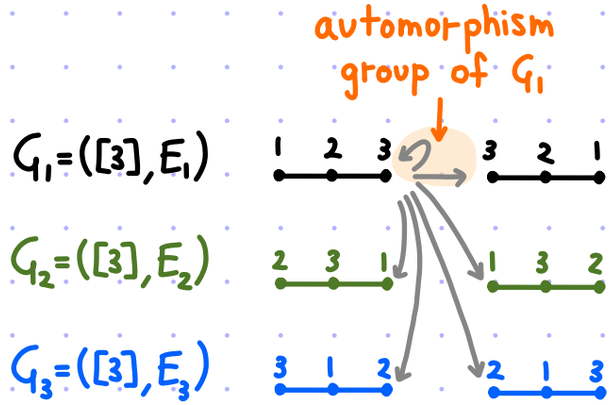

claim: G has $\frac{n!}{|aut(G)|}$ isomorphic graphs.

Hence $|\{(H,\pi) \mid H \equiv G \wedge \pi \in aut(H)\}| = n!$.

Given (G_0, G_1), define $S := \{(H,\pi) \mid (H \equiv G_0 \vee H \equiv G_1) \wedge \pi \in aut(H)\}$.

Observe that: $\begin{cases} G_0 \equiv G_1 \to |S| = n! \\ G_0 \not\equiv G_1 \to |S| = 2 \cdot n! \end{cases}$

Moreover, can prove that $(H,\pi) \in S$ by providing an isomorphism to G_0 or G_1.

→ It suffices for the prover to convince the verifier that $|S| \geq 2 \cdot n!$.

Tool: Pairwise Independent Hashing

NOTE:
$H_{m,\ell}$ is pairwise-independent
↳ $H_{m,\ell}$ is universal $\left(\forall x,x' \in \{0,1\}^m \text{ with } x \neq x',\ \Pr_{h \in H_{m,\ell}}[h(x)=h(x')] = \frac{1}{2^\ell}\right)$

A function family $H_{m,\ell} = \{ h: \{0,1\}^m \to \{0,1\}^\ell \}$ is **pairwise-independent** if

$$\forall\ x,x' \in \{0,1\}^m \text{ with } x \neq x',\ \forall\ y,y' \in \{0,1\}^\ell \quad \Pr_{h \in H_{m,\ell}}\begin{bmatrix} h(x)=y \\ h(x')=y' \end{bmatrix} = \frac{1}{2^{2\ell}}.$$

EXAMPLE: $H_{m,m} = \{ h_{a,b}(x) = ax+b \}_{a,b \in \mathbb{F}_{2^m}}$ (a random affine function over $\mathbb{F}_{2^m}$).

Indeed: $\Pr_{a,b}\begin{bmatrix} h_{a,b}(x)=y \\ h_{a,b}(x')=y' \end{bmatrix} = \Pr_{a,b}\begin{bmatrix} ax+b=y \\ ax'+b=y' \end{bmatrix} = \Pr_{a,b}\begin{bmatrix} a=\frac{y-y'}{x-x'} \\ b=y-ax \end{bmatrix} = \frac{1}{2^{2m}}.$

Actually we are interested in a family $H_{m,\ell}$ with $\ell < m$.

So consider: $H_{m,\ell} = \{ h_{a,b}(x) = ax+b \mod 2^\ell \}_{a,b \in \mathbb{F}_{2^m}}$.

The truncation to ℓ bits does NOT affect pairwise independence:
there are $2^{m-\ell}$ choices of $a \in \mathbb{F}_{2^m}$ s.t. $a\cdot(x-x') \mod 2^\ell = y-y'$,
and for each such $a \in \mathbb{F}_{2^m}$ there are $2^{m-\ell}$ choices of $b \in \mathbb{F}_{2^m}$ s.t. $ax+b \mod 2^\ell = y$.

We have an efficient pairwise-independent function family $H_{m,\ell}$ for every m,ℓ with $\ell \leq m$.

Set Lower Bound Protocol [1/2]

Let $S \subseteq \{0,1\}^m$ be in NP (i.e., can efficiently check that $x \in S$ with the help of a proof).

GOAL: an IP for the promise problem $\begin{cases} \text{YES if } |S| \geq B \\ \text{NO if } |S| \leq B/2 \end{cases}$.

P_S		$V_S(B)$
		Set $\ell \in \mathbb{N}$ s.t. $2^{\ell-2} \leq B \leq 2^{\ell-1}$.
Find $x \in S$ s.t. $h(x)=y$.	$\xleftarrow{h,y}$	Sample $h \leftarrow H_{m,\ell}$ and $y \leftarrow \{0,1\}^\ell$.
Find proof π for "$x \in S$".	$\xrightarrow{x,\pi}$	Check that $h(x)=y$ and π proves "$x \in S$".

lemma: if $|S| \geq B$ then $\Pr[\langle P_S, V_S(B)\rangle = 1] \geq \frac{3}{4}B \cdot \frac{1}{2^\ell}$

if $|S| \leq \frac{B}{2}$ then $\forall \tilde{P}\ \Pr[\langle \tilde{P}, V_S(B)\rangle = 1] \leq \frac{1}{2}B \cdot \frac{1}{2^\ell}$

} gap is $\geq \frac{1}{4}B \cdot \frac{1}{2^\ell} \geq \frac{1}{16}$.

Soundness: if $|S| \leq \frac{B}{2}$ then

$$\forall \tilde{P}\ \Pr[\langle \tilde{P}, V_S(B)\rangle = 1] \leq \Pr_{h,y}[\exists x \in S : h(x)=y] \leq \sum_{x \in S} \Pr_{h,y}[h(x)=y] \leq |S| \cdot \frac{1}{2^\ell} \leq \frac{1}{2}B \cdot \frac{1}{2^\ell}.$$

Set Lower Bound Protocol [2/2]

Let $S \subseteq \{0,1\}^m$ be in NP (i.e., can efficiently check that $x \in S$ with the help of a proof).

GOAL: an IP for the promise problem $\begin{cases} \text{YES if } |S| \geq B \\ \text{NO if } |S| \leq B/2 \end{cases}$.

P_S		$V_S(B)$
		Set $\ell \in \mathbb{N}$ s.t. $2^{\ell-2} \leq B \leq 2^{\ell-1}$.
Find $x \in S$ s.t. $h(x)=y$.	$\xleftarrow{h,y}$	Sample $h \leftarrow H_{m,\ell}$ and $y \leftarrow \{0,1\}^\ell$.
Find proof π for "$x \in S$".	$\xrightarrow{x,\pi}$	Check that $h(x)=y$ and π proves "$x \in S$".

lemma: if $|S| \geq B$ then $\Pr[\langle P_S, V_S(B)\rangle = 1] \geq \frac{3}{4}B \cdot \frac{1}{2^\ell}$

if $|S| \leq \frac{B}{2}$ then $\forall \tilde{P}\ \Pr[\langle \tilde{P}, V_S(B)\rangle = 1] \leq \frac{1}{2}B \cdot \frac{1}{2^\ell}$

} gap is $\geq \frac{1}{4}B \cdot \frac{1}{2^\ell} \geq \frac{1}{16}$.

Completeness: WLOG $|S|=B$ (larger $|S|$ increases acceptance probability). For every $y \in \{0,1\}^\ell$, (randomness of y is not used for completeness)

$$\Pr[\langle P_S, V_S(B)\rangle = 1] = \Pr_h[\exists x \in S : h(x)=y] \geq \sum_{x \in S} \Pr[h(x)=y] - \sum_{\substack{x,x' \in S \\ x \neq x'}} \Pr\begin{bmatrix} h(x)=y \\ h(x')=y \end{bmatrix} = |S| \cdot \frac{1}{2^\ell} - \binom{|S|}{2} \cdot \frac{1}{2^{2\ell}}$$

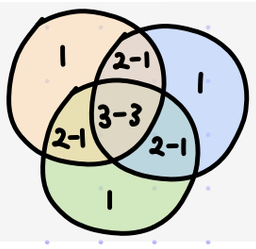

Inclusion-Exclusion Bound
$\Pr[\cup_i E_i] \geq \sum_i \Pr[E_i] - \sum_{i \neq j} \Pr[E_i \cap E_j]$

$$= B \cdot \frac{1}{2^\ell} - \binom{B}{2} \cdot \frac{1}{2^{2\ell}} \geq \frac{B}{2^\ell} - \frac{B^2}{2^{2\ell+1}} = \frac{B}{2^\ell} \cdot \left(1 - \frac{B}{2^{\ell+1}}\right) \geq \frac{B}{2^\ell} \cdot \left(1 - \frac{1}{4}\right) = \frac{3}{4}B \cdot \frac{1}{2^\ell}.$$

Public Coin Interactive Proof for GNI

theorem: $GNI \in AM[k=1]$

Apply the set lower bound protocol on $S := \left\{ (H,\pi) \in \{0,1\}^{n^2+n\log n} \;\middle|\; \begin{array}{l} (H \equiv G_0 \vee H \equiv G_1) \\ \wedge\ \pi \in aut(H) \end{array} \right\}$.

$P((G_0,G_1))$		$V((G_0,G_1))$
		$B := 2 \cdot n!$, $m := n^2 + n \cdot \log n$
		Set ℓ s.t. $2^{\ell-2} \leq B \leq 2^{\ell-1}$ [and so $\ell = O(n \cdot \log n)$]
Find $(H,\pi) \in S$ s.t. $h(H,\pi) = y$.	$\xleftarrow{h,\,y}$	Sample $h \leftarrow H_{m,\ell}$ and $y \leftarrow \{0,1\}^{\ell}$.
Find isomorphism φ from H to G_b.	$\xrightarrow{(H,\pi),\,\varphi}$	Check that $h(H,\pi) = y$ and $(H,\pi) \in S$. [$(\varphi(H) = G_0 \vee \varphi(H) = G_1) \wedge \pi \in aut(H)$]

Completeness: if $(G_0,G_1) \in GNI$ then $|S| = 2 \cdot n!$ so

$$\Pr\left[\langle P((G_0,G_1)), V((G_0,G_1))\rangle = 1\right] = \Pr_{h,y}\left[\exists (H,\pi) \in S : h(H,\pi) = y\right] \geq \frac{3}{4} \cdot \frac{B}{2^{\ell}}.$$

Soundness: if $(G_0,G_1) \notin GNI$ then $|S| = n!$ so

$$\forall \tilde{P}\ \Pr\left[\langle \tilde{P}, V((G_0,G_1))\rangle = 1\right] = \Pr_{h,y}\left[\exists (H,\pi) \in S : h(H,\pi) = y\right] \leq \frac{1}{2} \cdot \frac{B}{2^{\ell}}.$$

Perfect Completeness for Public Coins

The set lower bound protocol introduces a completeness error.

This is NOT essential:

On Completeness and Soundness in Interactive Proof Systems

Martin Fürer	Oded Goldreich	Yishay Mansour	Michael Sipser	Stathis Zachos
Penn State	Technion	MIT	MIT	CUNY





theorem:

If L has a k-round public-coin IP

then L has a $(k+1)$-round public-coin IP with perfect completeness.

Example: We showed that $GNI \in AM[k=1]$, so we deduce that $GNI \in AM[\varepsilon_c = 0, k=2]$.

(GNI has a 2-round public-coin IP with perfect completeness.)

We proceed in several steps.

- Warmup: simple protocol to reduce (but not eliminate) completeness error.
- Review: Lautemann's proof that $BPP \subseteq \Sigma_2^P$.
- Proof: we build on warmup and review.

Warmup: Reduce Completeness Error

Repeat the protocol multiple times and accept if AT LEAST ONE execution accepts.

$P_*(x)$:

$\forall i \in [t],\ a_j^{(i)} := P(x, \rho_1^{(i)}, \dots, \rho_{j-1}^{(i)})$

For $j=1,\dots,k$:

$a_j^{(1)}, \dots, a_j^{(t)} \rightarrow$

$\leftarrow \rho_j^{(1)}, \dots, \rho_j^{(t)}$

$V_*(x)$:

Sample $\rho_j^{(1)}, \dots, \rho_j^{(t)} \in \{0,1\}^{r_j}$.

$\exists i \in [t]\ V(x, a_1^{(i)}, \dots, a_k^{(i)}; \rho_j) = 1$

For every repetition parameter $t \in \mathbb{N}$:

- $\varepsilon_c \mapsto \varepsilon_c' = \varepsilon_c^t$ — $\Pr[\langle P_*(x), V_*(x)\rangle = 0] = (\Pr[\langle P(x), V(x)\rangle = 0])^t \leq \varepsilon_c^t$
- $\varepsilon_s \mapsto \varepsilon_s' = t \cdot \varepsilon_s$ — $\Pr[\langle P_*(x), V_*(x)\rangle = 1] \leq t \cdot \Pr[\langle P(x), V(x)\rangle = 1] \leq t \cdot \varepsilon_s$
- $k \mapsto k' = k$ — The t executions are in parallel.
- $c \mapsto c' = t \cdot c$ — Each execution contributes c bits of communication.

The completeness error can be made arbitrarily small, but NOT zero.

BUT: a clever twist on this protocol achieves perfect completeness.

Review: Lautemann Theorem [1/2]

theorem: $BPP \subseteq \Sigma_2^P$ Review: $L \in \Sigma_2^P \leftrightarrow \exists$ polynomial-time algorithm D s.t. $\begin{cases} x \in L \rightarrow & \exists y \; \forall z \; D(x,y,z)=1 \\ x \notin L \rightarrow & \forall y \; \exists z \; D(x,y,z)=0 \end{cases}$.

Let L be decidable by a polynomial-time probabilistic algorithm M with $\begin{cases} \text{YES-error probability} \leq \alpha \\ \text{NO-error probability} \leq \beta \end{cases}$.

PROBLEM: The probabilistic algorithm M sometimes errs.

How to construct a deterministic algorithm D that NEVER errs?

We have to somehow "get rid" of errors caused by randomness.

IDEA: Consider multiple correlated executions of M via randomness shifts.

- $x \in L \rightarrow \Pr[M(x)=0] \leq \alpha$

  Many accepting randomness strings.

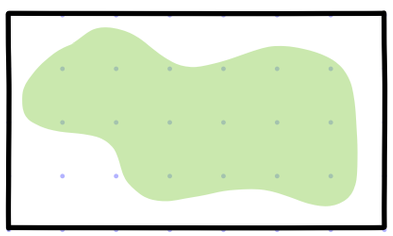

$\forall \rho \; \exists i \in [t]$ $\rho \oplus \sigma^{(i)}$ is accepting

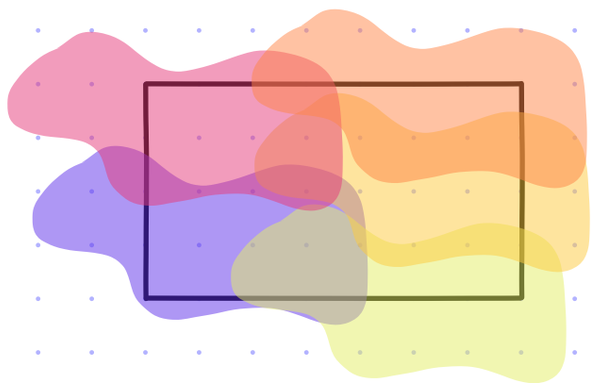

We can find a small number of shifts $\sigma^{(1)}, \ldots, \sigma^{(t)}$ such that ...

- $x \notin L \rightarrow \Pr[M(x)=1] \leq \beta$

  Few accepting randomness strings.

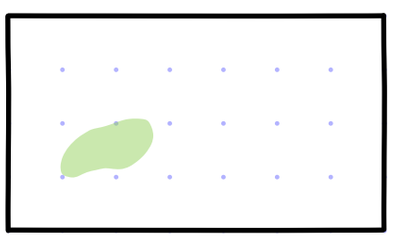

$\exists \rho \; \forall i \in [t]$ $\rho \oplus \sigma^{(i)}$ is rejecting

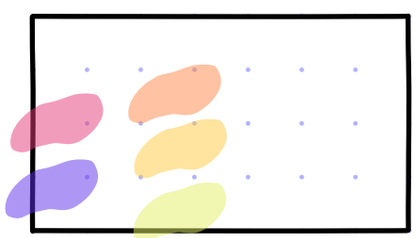

Review: Lautemann Theorem [2/2]

theorem: $BPP \subseteq \Sigma_2^P$ Review: $L \in \Sigma_2^P \leftrightarrow \exists$ polynomial-time algorithm D s.t. $\begin{cases} x \in L \to \exists y\, \forall z\, D(x,y,z)=1 \\ x \notin L \to \forall y\, \exists z\, D(x,y,z)=0 \end{cases}$.

Let L be decidable by a polynomial-time probabilistic algorithm M with $\begin{cases} \text{YES-error probability} \le \alpha \\ \text{NO-error probability} \le \beta \end{cases}$.

We use the probabilistic method to show the two conditions.

- If $x \in L$ then (provided $t > -\frac{r}{\log \alpha}$) $\exists \sigma^{(1)},\dots,\sigma^{(t)} \in \{0,1\}^r$ $\forall \varsigma \in \{0,1\}^r$ $\left(\exists\, i \in [t]\ M(x;\sigma^{(i)} \oplus \varsigma)=1\right)$:

$$\Pr_{\sigma^{(1)},\dots,\sigma^{(t)}}\left[\exists \varsigma \in \{0,1\}^r \left(\forall\, i \in [t]\ M(x;\sigma^{(i)} \oplus \varsigma)=0\right)\right] \le \sum_{\varsigma \in \{0,1\}^r} \Pr_{\sigma^{(1)},\dots,\sigma^{(t)}}\left[\forall\, i \in [t]\ M(x;\sigma^{(i)} \oplus \varsigma)=0\right]$$

For t large enough MOST $\sigma^{(1)},\dots,\sigma^{(t)}$ are good.

$$= 2^r \cdot \Pr_{\varsigma^{(1)},\dots,\varsigma^{(t)}}\left[\forall\, i \in [t]\ M(x;\varsigma^{(i)})=0\right] \le 2^r \cdot \alpha^t < 1.$$

- If $x \notin L$ then (provided $t < \frac{1}{\beta}$) $\forall \sigma^{(1)},\dots,\sigma^{(t)} \in \{0,1\}^r$ $\exists \varsigma \in \{0,1\}^r$ $\left(\forall\, i \in [t]\ M(x;\sigma^{(i)} \oplus \varsigma)=0\right)$:

  Fix $\sigma^{(1)},\dots,\sigma^{(t)} \in \{0,1\}^r$. For every $i \in [t]$, $\Pr_{\varsigma \in \{0,1\}^r}\left[M(x;\sigma^{(i)} \oplus \varsigma)=1\right] = \Pr_{\varsigma \in \{0,1\}^r}\left[M(x;\varsigma)=1\right] \le \beta$.

  Hence, $\Pr_{\varsigma \in \{0,1\}^r}\left[\exists\, i \in [t]\ M(x;\sigma^{(i)} \oplus \varsigma)=1\right] \le \sum_{i \in [t]} \Pr_{\varsigma \in \{0,1\}^r}\left[M(x;\sigma^{(i)} \oplus \varsigma)=1\right] \le t \cdot \beta < 1.$

The condition $\exists t \in \mathbb{N}\ -\frac{r}{\log \alpha} < t < \frac{1}{\beta}$ can be achieved by repetition (and taking majority).

E.g. for $\alpha_0, \beta_0 = 1/3$: ℓ-wise error reduction gives $r = \ell \cdot r_0$ and $\alpha, \beta = \exp(-\ell)$, yielding $\mathrm{poly}(\ell) \cdot r_0 < t < \exp(\ell)$.

Proof of Perfect Completeness for IPs [1/3]

Let (P,V) be a k-round public-coin IP for L.

Let r be V's randomness complexity, divided by rounds as $r_1,\dots,r_k$ with $\sum_{j\in[k]} r_j = r$.

For every repetition parameter $t\in\mathbb{N}$ the new public-coin IP (P_*,V_*) is as follows.

$P_*(x)$

Find $\sigma^{(1)},\dots,\sigma^{(t)}\in\{0,1\}^r$ s.t.

$\forall \rho\in\{0,1\}^r \ \exists i\in[t] \ \langle P(x),V(x,\sigma^{(i)}\oplus\rho)\rangle = 1$

$\forall i\in[t],\ a_j^{(i)} := P(x,\sigma_1^{(i)}\oplus\rho_1,\dots,\sigma_{j-1}^{(i)}\oplus\rho_{j-1})$

$\xrightarrow{\sigma^{(1)},\dots,\sigma^{(t)}}$

$\xrightarrow{a_j^{(1)},\dots,a_j^{(t)}}$

$\xleftarrow{\rho_j}$

$V_*(x)$

For $j=1,\dots,k$:

Sample $\rho_j\in\{0,1\}^{r_j}$.

$\exists i\in[t] \ V(x,a_1^{(i)},\dots,a_k^{(i)};\sigma^{(i)}\oplus\rho)=1$.

- $\varepsilon_c \mapsto \varepsilon_c' = 0$ Provided that $t > -\frac{r}{\log\varepsilon_c}$, as we prove soon.
- $\varepsilon_s \mapsto \varepsilon_s' = t\cdot\varepsilon_s$ As we prove soon. It is <1 provided that $t < \frac{1}{\varepsilon_s}$.
- $k \mapsto k' = k+1$ There are t (correlated) executions in parallel, plus an extra message.
- $c \mapsto c' = t\cdot(c+r)$ Each execution contributes c bits, plus $t\cdot r$ bits in the extra message.

The condition $\exists t\in\mathbb{N} \ -\frac{r}{\log\varepsilon_c} < t < \frac{1}{\varepsilon_s}$ can be achieved by repetition (and taking majority).

Proof of Perfect Completeness for IPs [2/3]

$P_*(x)$

Find $\sigma^{(1)},\dots,\sigma^{(t)} \in \{0,1\}^r$ s.t.

$\forall \rho \in \{0,1\}^r \; \exists i \in [t] \; \langle P(x), V(x, \sigma^{(i)} \oplus \rho) \rangle = 1$

$\sigma^{(1)},\dots,\sigma^{(t)}$ →

$\forall i \in [t]$, $a_j^{(i)} := P(x, \sigma_1^{(i)} \oplus \rho_1, \dots, \sigma_{j-1}^{(i)} \oplus \rho_{j-1})$

For $j = 1,\dots,k$:

$a_j^{(1)},\dots,a_j^{(t)}$ →

← ρ_j

$V_*(x)$

Sample $\rho_j \in \{0,1\}^{r_j}$.

$\exists i \in [t] \; V(x, a_1^{(i)},\dots,a_k^{(i)}; \sigma^{(i)} \oplus \rho) = 1$.

Completeness:

Suppose that $x \in L$.

If $P_*(x)$ finds "good" $\sigma^{(1)},\dots,\sigma^{(t)}$ ($\forall \rho \in \{0,1\}^r \; (\exists i \in [t] \; \langle P(x), V(x, \sigma^{(i)} \oplus \rho) \rangle = 1)$) then $P_*(x)$ convinces $V_*(x)$ with probability 1.

They exist:

$$\Pr_{\sigma^{(1)},\dots,\sigma^{(t)}}\left[\exists \rho \in \{0,1\}^r \; \forall i \in [t] \; \langle P(x), V(x, \sigma^{(i)} \oplus \rho) \rangle = 0\right] \le \sum_{\rho \in \{0,1\}^r} \Pr_{\sigma^{(1)},\dots,\sigma^{(t)}}\left[\forall i \in [t] \; \langle P(x), V(x, \sigma^{(i)} \oplus \rho) \rangle = 0\right]$$

$$= 2^r \cdot \Pr_{\rho^{(1)},\dots,\rho^{(t)}}\left[\forall i \in [t] \; \langle P(x), V(x, \rho^{(i)}) \rangle = 0\right] \le 2^r \cdot \varepsilon_c^t < 1.$$

$t > -\frac{r}{\log \varepsilon_c}$

Proof of Perfect Completeness for IPs [3/3]

Soundness: Suppose that $x \notin L$ and fix a malicious prover $\tilde{P}_*$.

For every $i \in [t]$, define $\tilde{P}_i$ against V as follows:

- Run $\tilde{P}_*$ to obtain $(\sigma^{(1)}, \ldots, \sigma^{(t)})$.
- In round $j \in [k]$ (having received $q_1, \ldots, q_{j-1}$ from V):
  - compute the next message as $a_j := \tilde{P}_*(q_1 \oplus \sigma_1^{(i)}, \ldots, q_{j-1} \oplus \sigma_{j-1}^{(i)})[i]$.

Define $(\sigma^{(1)}, \ldots, \sigma^{(t)}) := \tilde{P}_*$ (the prover's first message).

For every $i \in [t]$,

$$\Pr_{q \in \{0,1\}^r}\left[V(x, \tilde{P}_*(q_1)[i], \ldots, \tilde{P}_*(q_1, \ldots, q_k)[i]; \sigma^{(i)} \oplus q) = 1 \right]$$

$$= \Pr_{q \in \{0,1\}^r}\left[V(x, \tilde{P}_i(\sigma_1^{(i)} \oplus q_1), \ldots, \tilde{P}_i(\sigma_1^{(i)} \oplus q_1, \ldots, \sigma_k^{(i)} \oplus q_k); \sigma^{(i)} \oplus q) = 1 \right]$$

$$= \Pr_{q \in \{0,1\}^r}\left[V(x, \tilde{P}_i(q_1), \ldots, \tilde{P}_i(q_1, \ldots, q_k); q) = 1 \right] \leq \varepsilon_s .$$

We conclude that

$$\Pr_{q \in \{0,1\}^r}\left[\langle \tilde{P}_*, V_*(x; q) \rangle = 1 \right] = \Pr_{q \in \{0,1\}^r}\left[\exists\, i \in [t] \;\; V(x, \tilde{P}_*(q_1)[i], \ldots, \tilde{P}_*(q_1, \ldots, q_k)[i]; \sigma^{(i)} \oplus q) = 1 \right]$$

$$\leq \sum_{i \in [t]} \Pr_{q \in \{0,1\}^r}\left[V(x, \tilde{P}_*(q_1)[i], \ldots, \tilde{P}_*(q_1, \ldots, q_k)[i]; \sigma^{(i)} \oplus q) = 1 \right] \leq t \cdot \varepsilon_s < 1.$$

$t < \frac{1}{\varepsilon_s}$

The Case of IOPs: Private to Public Coins

The IP transformation to obtain a public-coin verifier can be applied to an IOP (view it as an IP)
but the resulting IOP has bad query complexity.
Indeed, the set lower bound protocol requires the verifier to read the (alleged) preimage $x \in S$
and the size of x would be at least the size of a prover message (an IOP string in this case).

Nevertheless a similar IOP-friendly transformation exists:

theorem: If $L \in \mathrm{IOP}[\varepsilon_c, \varepsilon_s, k, \Sigma, \ell, q, r]$
then $L \in \mathrm{IOP}[\text{public-coin}, \varepsilon_c' < 1, \varepsilon_s' < 1, k' = O(k), \Sigma = \{0,1\}, \ell' = \mathrm{poly}(\ell), q = \text{"}O(1) \text{ per round"}, r' = r + O(\log n)]$

The proof approach is as follows:

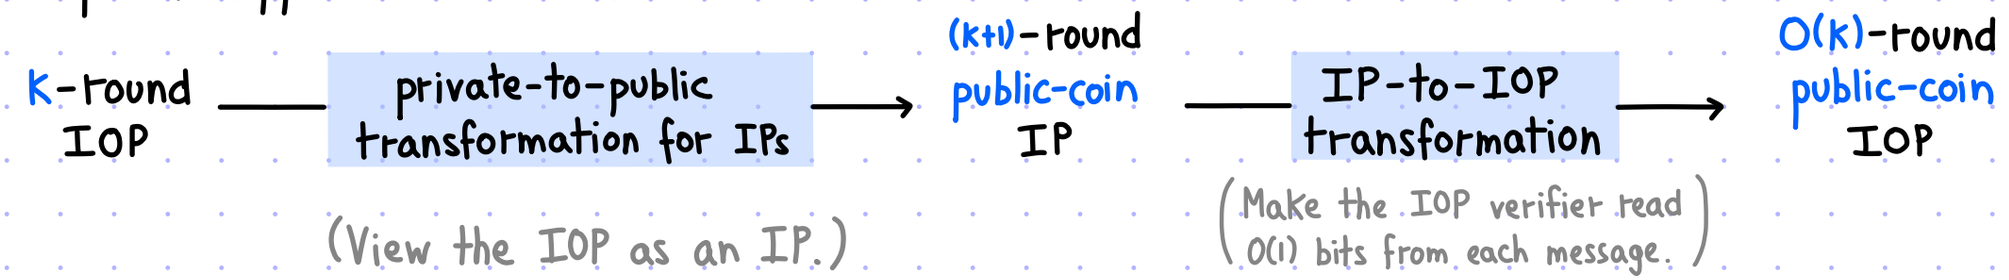

A key ingredient of the IP-to-IOP transformation is INDEX-DECODABLE PCPs,
a strengthening of the notion of Holographic PCPs.

REMARK: A PCP is a public-coin IOP, so the public-vs-private-coin question is trivial.

The Case of IOPs: Perfect Completeness

The IP transformation for perfect completeness extends to IOPs, with a moderate increase in query complexity: $q \mapsto q' = t \cdot q = O\left(-\frac{r}{\log \varepsilon_c}\right) \cdot q$.

PROBLEM: typically $r = \Omega(\log n)$, so q' is super-constant even if q is constant.

We can preserve query complexity (up to a small additive constant) with a small tweak.

$P_*(x)$		$V_*(x)$
Find $\sigma^{(1)}, \dots, \sigma^{(t)} \in \{0,1\}^r$ s.t. $\forall \rho \in \{0,1\}^r \; \exists i \in [t] \; \langle P(x), V(x, \sigma^{(i)} \oplus \rho) \rangle = 1$	$\xrightarrow{\sigma^{(1)}, \dots, \sigma^{(t)}}$	
$\forall i \in [t], a_j^{(i)} := P(x, \sigma_1^{(i)} \oplus \rho_1, \dots, \sigma_{j-1}^{(i)} \oplus \rho_{j-1})$	For $j = 1, \dots, k$: $\xrightarrow{a_j^{(1)}, \dots, a_j^{(t)}}$ $\xleftarrow{\rho_j}$	Sample $\rho_j \in \{0,1\}^{r_j}$.
Find $i \in [t]$ s.t. $\langle P(x), V(x, \sigma^{(i)} \oplus \rho) \rangle = 1$	$\xrightarrow{i}$	$V(x, a_1^{(i)}, \dots, a_k^{(i)}; \sigma^{(i)} \oplus \rho) = 1$.

The IOP prover tells the IOP verifier which execution accepts.

➡ The IOP verifier reads $i \in [t]$, then reads $\sigma^i \in \{0,1\}^r$, and then checks the i-th execution with randomness $\sigma^{(i)} \oplus \rho$.

NOTE: the IOP verifier is adaptive.

New parameters:

- $\varepsilon_c \mapsto \varepsilon_c' = 0$
- $\varepsilon_s \mapsto \varepsilon_s' = t \cdot \varepsilon_s$
- $k \mapsto k' = k+1$
- $|\Sigma| \mapsto |\Sigma'| = \max\{|\Sigma|, 2^r, t\}$
- $\ell \mapsto \ell' = t \cdot \ell + t + 1$
- $q \mapsto q' = q+2$
- $r \mapsto r' = r$

Bibliography

Private-coin to public-coin

- [GS 1986]: Private coins versus public coins in interactive proof systems, by Shafi Goldwasser, Michael Sipser.
- [GL 2016]: On emulating interactive proofs with public coins, by Oded Goldreich, Maya Leshkowitz.
- [AR 2021]: On prover-efficient public-coin emulation of interactive proofs, by Gal Arnon, Guy Rothblum.
- [ACY 2022]: A PCP theorem for interactive proofs and applications, by Gal Arnon, Alessandro Chiesa, Eylon Yogev.

Perfect completeness

- [FGMSZ 1989]: On completeness and soundness in interactive proof systems, by Martin Furer, Oded Goldreich, Yishay Mansour, Michael Sipser, Stathis Zachos.
- [ABCY 2022]: A toolbox for barriers on interactive oracle proofs, by Gal Arnon, Amey Bhangale, Alessandro Chiesa, Eylon Yogev.
- [ABY 2024]: Hamming weight proofs of proximity with one-sided error, by Gal Arnon, Shany Ben-David, Eylon Yogev.