Lecture B.9

Proof Composition

Summer Graduate School on
Foundations and Frontiers of Probabilistic Proofs
2021.08.05

Proof Composition

We have seen techniques to achieve either

(i) polynomial proof length and polylogarithmic query complexity, OR

(ii) exponential proof length and constant query complexity

How to achieve the best of both?

Recap: Reed—Muller PCP

NP-complete => for all NP

We showed a poly-length, polylog-local PCP for QESAT.

The PCP was a Reed-Muller enc. of sat. assignment $a_1, \dots, a_n$

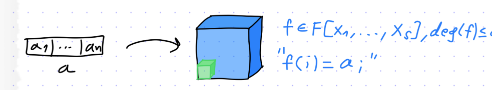

Π_{sc} is eval table of IP prover for sumcheck claim $\sum_{\alpha, \beta \in H_v^{s_v}} g(\alpha, \beta) = 0$

We used the structure of the Reed-Muller code to:

① Reduce m equations to 1 using pseudo-random linear comb.

Recap: Hadamard PCP

NP-complete => for all NP

We showed an exp-length, $O(1)$-local PCP for QESAT.

The PCP had 2 parts: (I) Hadamard enc. of sat. assignment $a_1, \dots, a_n$

(II) Hadamard enc. of $a \otimes a = (a_i \cdot a_j)_{i,j}$

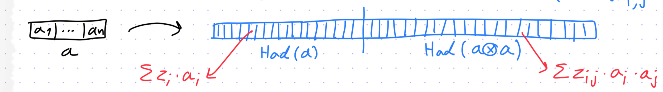

We used the linear structure of the Hadamard code to:

① check all equation at once using random linear comb.

We will learn about PROOF COMPOSITION: a technique to combine two PCPs so that the composed PCP inherits the proof length of one PCP and the query complexity of the other PCP.

In particular this leads to a result known as the PCP Theorem:

$$NP \subseteq PCP[\varepsilon_c = 0, \varepsilon_s = 1/2, \Sigma = \{0,1\}, \ell = poly(n), q = O(1), r = O(\log n)].$$

We will also learn about INTERACTIVE PROOF COMPOSITION, which works for IOPs.

For example, this leads to an optimal tradeoff between proof length & query complexity:

$$CSAT \subseteq IOP[\varepsilon_c = 0, \varepsilon_s = 1/2, k = 3, \Sigma = \{0,1\}, \ell = O(n), q = O(1), r = O(\log n)].$$

High-Level Plan

Ingredients: (i) an outer PCP (P_{out}, V_{out}) for a language L — "good" proof length

(ii) an inner PCP (P_{in}, V_{in}) for the relation $R(V_{out})$ — "good" query complexity

We wish to construct a new PCP (P,V) for the language L with the best of both.

Idea: use the inner PCP to check the computation of the outer PCP's verifier

[this is reminiscent of code concatenation in coding theory for reducing alphabet size]

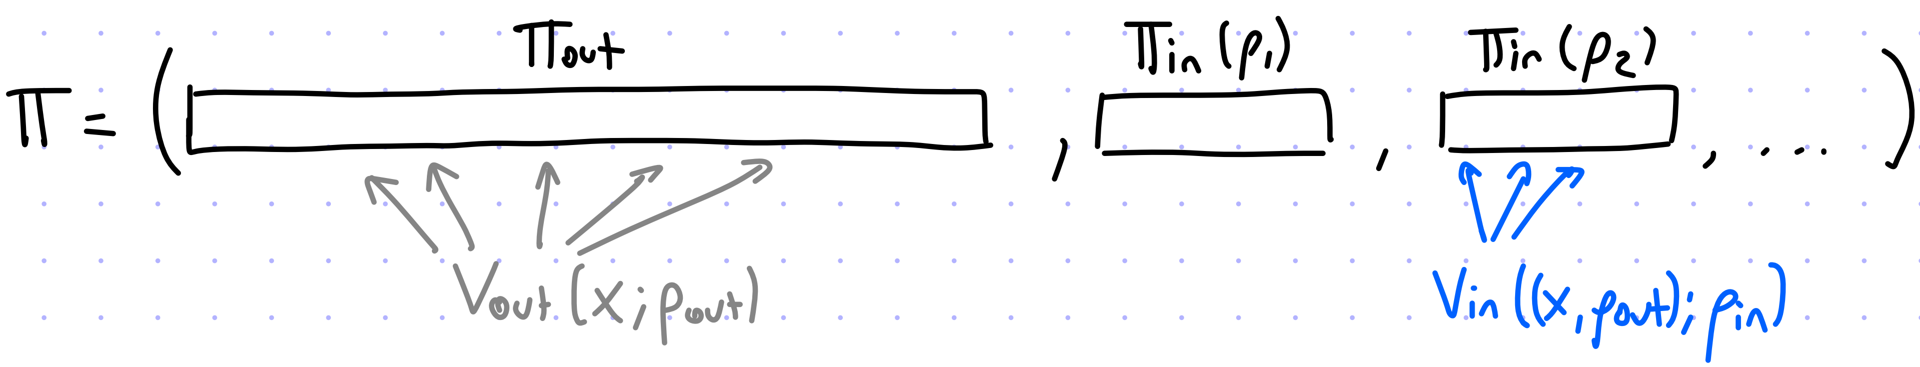

$P(x)$

1. Compute outer PCP: $\pi_{out} := P_{out}(x)$
2. For each $\rho_{out} \in \{0,1\}^{r_{out}}$:

   compute inner PCP for ρ_{out} as

   $\pi_{in}[\rho_{out}] := P_{in}((x, \rho_{out}))$
3. Output $\pi := (\pi_{out}, (\pi_{in}[\rho_{out}])_{\rho_{out} \in \{0,1\}^{r_{out}}})$.

$V^{\pi}(x)$

1. Sample $\rho_{out} \in \{0,1\}^{r_{out}}$.
2. Check that $V_{in}^{\pi_{in}[\rho_{out}]}(\underbrace{(x, \rho_{out})}_{x_{in}}) = 1$.

This plan has some problems...

Problems with the Plan

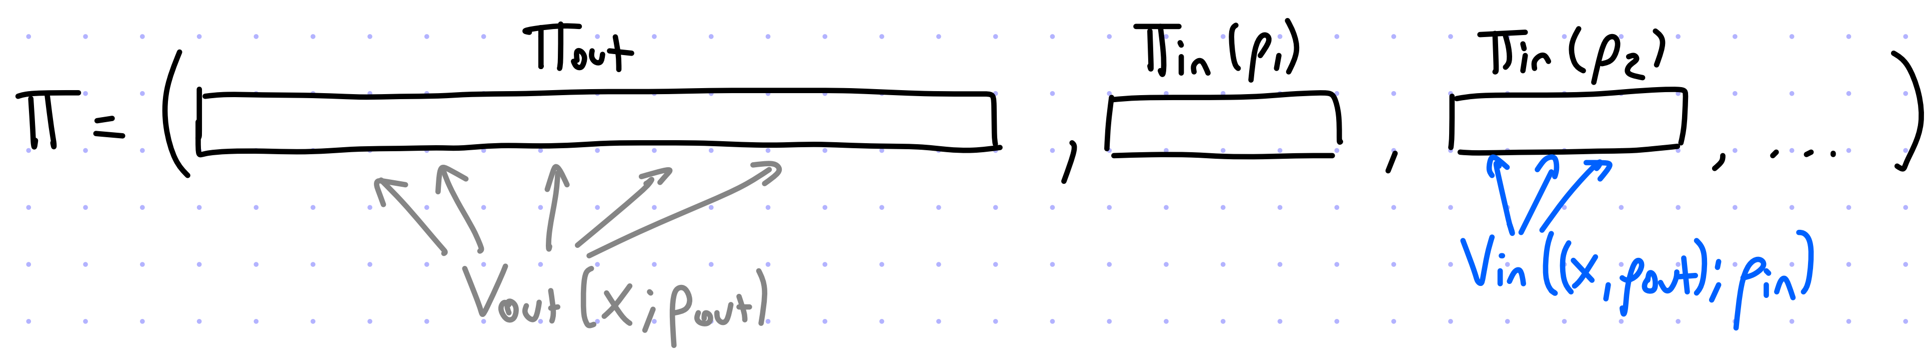

- **Problem:** To reduce query complexity we need to locally check both proof and statement!

Approach: Each inner PCP should be a "proof of proximity" for the corresponding local view. I.e., property testing "is this local view (derived from the given Π_{out}) satisfying for (x, ρ_{out})?"

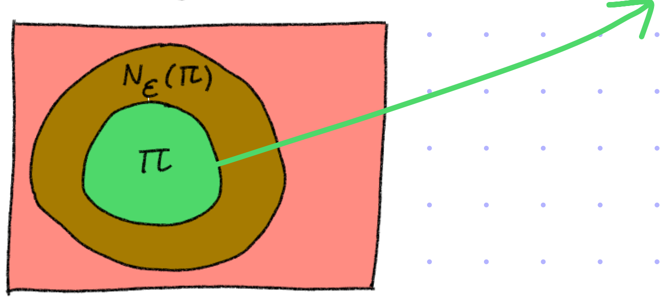

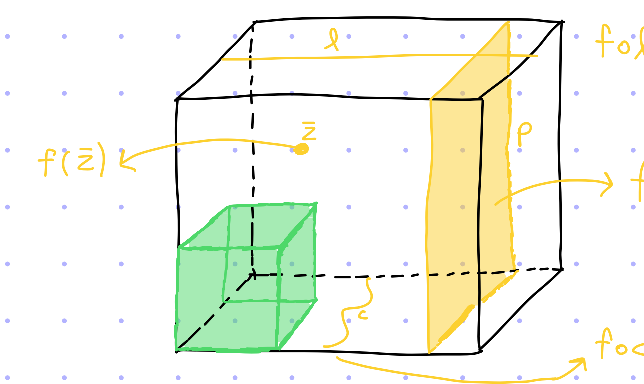

- **Problem:** We cannot hope to detect with a small number of queries to a local view whether the local view is accepting or rejecting. (Maybe it differs in 1 location from an accepting one!)

Approach: The outer PCP should be robust, i.e., if $x \notin L$ then whp a local view is far from any accepting local view.

Robust PCPs [for outer PCP]

We restrict attention on non-adaptive verifiers, which can be viewed as follows:

$V^{\pi}(x;\rho) = D(x, \pi[Q(x,\rho)], \rho)$ where $\begin{cases} Q \text{ is the query algorithm of } V \\ D \text{ is the decision algorithm of } V \end{cases}$

This induces the relation of accepting local views for the verifier V:

$R(V) := \{((x,\rho),a) \mid a \in \Sigma^{Q(x,\rho)} \wedge D(x,a,\rho)=1\}$

<u>def:</u> (P,V) is a PCP system for a language L with robustness parameter σ if:

① <u>completeness:</u> $\forall x \in L$, for $\pi := P(x)$, $\Pr_{\rho}[V^{\pi}(x;\rho)=1] \geq 1-\varepsilon_c$

accepting local view for (x,ρ)
$\{a \mid ((x,\rho),a) \in R(V)\}$

② <u>robust soundness:</u> $\forall x \notin L \; \forall \tilde{\pi} \; \Pr_{\rho}[\Delta(\tilde{\pi}[Q(x,\rho)], R(V)[(x,\rho)]) \leq \sigma] \leq \varepsilon_s$

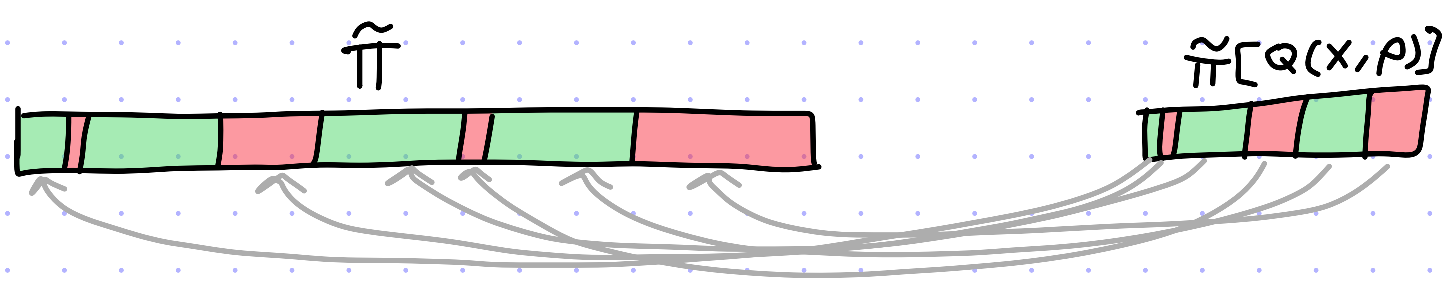

<u>Note:</u> Standard soundness is above definition with $\sigma = 0$: $V^{\tilde{\pi}}(x;\rho)=1 \leftrightarrow \Delta(\tilde{\pi}[Q(x,\rho)], R(V)[(x,\rho)])=0$.

PCPs of Proximity

A PCPP is to prove, for a given instance x and candidate witness w, that w is close to a valid witness for x (if one exists). The PCPP verifier has oracle access to w (and a proof).

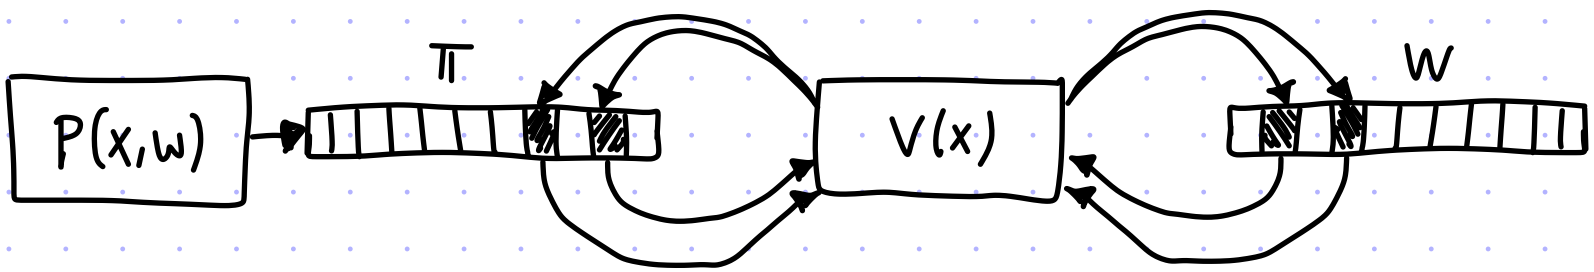

The query complexity counts queries to w & π.

Let $R = \{(x,w) \mid \dots\}$ be a binary relation.

Define
- the language of R: $L(R) = \{x \mid \exists w \text{ s.t. } (x,w) \in R\}$
- the valid witnesses of x: $R[x] = \{w \mid (x,w) \in R\}$ [if $x \notin L(R)$ then $R[x] = \emptyset$]

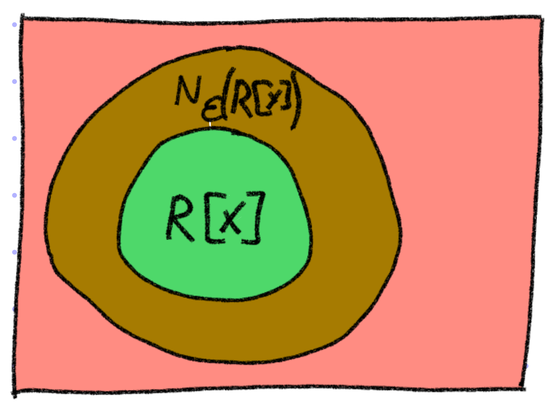

<u>def</u>: (P,V) is a PCPP system for a relation R with proximity parameter δ if:

① <u>completeness</u>: $\forall (x,w) \in R$, for $\pi := P(x,w)$, $\Pr_\rho[V^{w,\pi}(x;\rho)=1] \geq 1-\varepsilon_c$

[convention: $\Delta(w,\emptyset) := 1$]

② <u>proximity soundness</u>: $\forall (x,w)$ if $\Delta(w,R[x]) \geq \delta$ then $\forall \tilde{\pi}$ $\Pr_\rho[V^{w,\tilde{\pi}}(x;\rho)=1] \leq \varepsilon_s$

The Composed PCP

Ingredients: (i) outer: non-adaptive PCP (P_{out}, V_{out}) for a language L with robustness σ_{out}

(ii) inner: PCP of proximity (P_{in}, V_{in}) for the relation $R(V_{out})$ with proximity δ_{in}

The new PCP (P,V) for the language L is defined as follows:

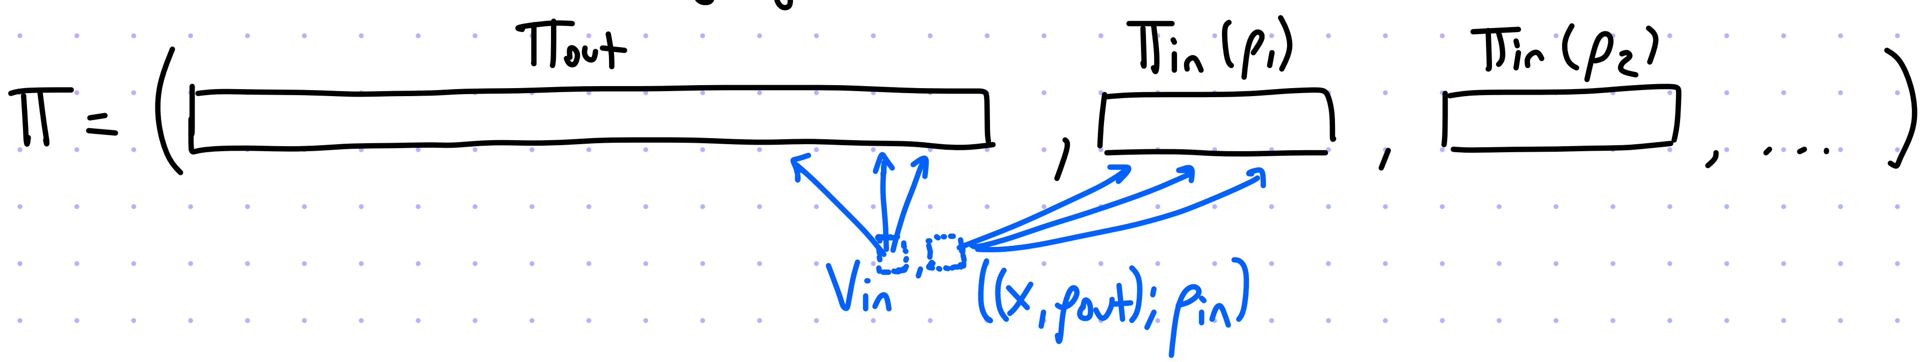

$P(x)$

1. Compute outer PCP: $\Pi_{out} := P_{out}(x)$
2. For each $\rho_{out} \in \{0,1\}^{r_{out}}$:

   compute inner PCPP for ρ_{out} as

   $\Pi_{in}[\rho_{out}] := P_{in}((x,\rho_{out}), \Pi_{out}[Q_{out}(x,\rho_{out})])$
3. Output $\Pi := (\Pi_{out}, (\Pi_{in}[\rho_{out}])_{\rho_{out}\in\{0,1\}^{r_{out}}})$.

$V^{\Pi}(x)$

1. Sample $\rho_{out} \in \{0,1\}^{r_{out}}$.
2. Check that $V_{in}^{\underbrace{\Pi_{out}[Q_{out}(x,\rho_{out})]}_{w_{in}}, \Pi_{in}[\rho_{out}]}(\underbrace{(x,\rho_{out})}_{x_{in}}) = 1$.

Soundness: If $x \notin L$, except w.p. ε_{out} over $\rho_{out} \in \{0,1\}^{r_{out}}$, the local view $\Pi_{out}[Q_{out}(x,\rho_{out})]$ is σ_{out}-far from $R(V_{out})[(x,\rho_{out})]$. If so (and $\sigma_{out} \geq \delta_{in}$) then V_{in} accepts w.p. $\leq \varepsilon_{in}$ over $\rho_{in} \in \{0,1\}^{r_{in}}$.
Overall soundness error is $\varepsilon = \varepsilon_{out} + \varepsilon_{in}$.

Proof Composition Theorem

Ingredients: (i) outer : a non-adaptive PCP (P_{out}, V_{out}) for a language L with robustness σ_{out}

(ii) inner : a PCP of proximity (P_{in}, V_{in}) for the relation $R(V_{out})$ with proximity δ_{in}

theorem: Then we get a PCP (P,V) for the language L s.t. if $\sigma_{out} \geq \delta_{in}$

- soundness error : $\varepsilon = \varepsilon_{out} + \varepsilon_{in}$
- randomness complexity: $r = r_{out} + r_{in}$
- proof length: $\ell = \ell_{out} + 2^{r_{out}} \cdot \ell_{in}$ (and similarly for prover time: $pt = pt_{out} + 2^{r_{out}} \cdot pt_{in}$)
- query complexity: $q = q_{in}$ (and similarly for verifier time: $vt = vt_{in}$)

How do we use it to prove the PCP theorem? [sketch]

(I) Observe the Hadamard PCP can be viewed as a PCPP.

(II) The Reed-Muller PCP can be made robust because RM is.

Problem: to get poly-length and $O(1)$-queries, we need to compose twice,

Solution: If both parts are robust-PCPP, so is the compose PCPP!

Proof Composition For IOPs?

We can similarly define robust IOPs and IOPs of proximity:

- def: (P,V) is an IOP system for a language L with robustness parameter σ if:

  ① completeness: $\forall x \in L \quad \Pr_{\rho}\left[\langle P(x), V(x;\rho)\rangle = 1\right] \geq 1-\varepsilon_c$

  accepting local view for (x,ρ)
  $\{a \mid ((x,\rho),a) \in R(V)\}$

  ② robust soundness: $\forall x \notin L \ \forall \tilde{P} \quad \Pr_{\rho}\left[\Delta(\tilde{\pi}[Q(x,\rho)], R(V)[(x,\rho)]) \leq \sigma \text{ where } \tilde{\pi} = \text{oracles}\left(\langle \tilde{P}, V(x;\rho)\rangle\right)\right] \leq \varepsilon_s$

- def: (P,V) is an IOPP system for a relation R with proximity parameter δ if:

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

  [convention: $\Delta(w,\emptyset) := 1$]

  ② proximity soundness: $\forall (x,w)$ if $\Delta(w, R(x)) \geq \delta$ then $\forall \tilde{P} \quad \Pr_{\rho}\left[\langle \tilde{P}, V^{w}(x;\rho)\rangle = 1\right] \leq \varepsilon_s$

Ex: if we set $R = \{((\mathbb{F},L,d),f) \mid f \in RS[\mathbb{F},L,d]\}$ then we get an IOPP for the Reed–Solomon code, of which FRI is an example.

Interactive Proof Composition

Ingredients: (i) outer: non-adaptive IOP (P_{out}, V_{out}) for a language L with robustness σ_{out}

(ii) inner: IOP of proximity (P_{in}, V_{in}) for the relation $R(V_{out})$ with proximity δ_{in}

For composition, the new IOP verifier tells the IOP prover which ρ_{out} it chose:

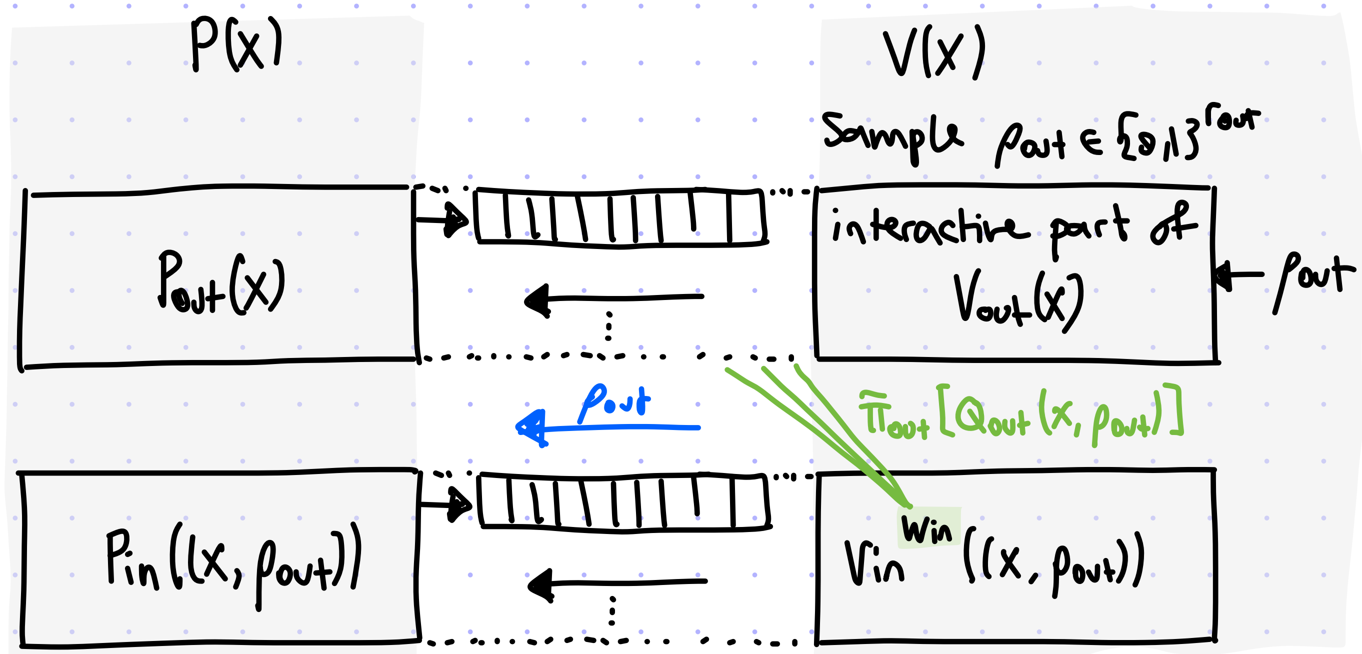

There is NO need to run inner IOP for every $\rho_{out} \in \{0,1\}^{r_{out}}$.

theorem: Then we get an IOP (P,V) for the language L s.t. if $\sigma_{out} \geq \delta_{in}$:

- soundness error: $\varepsilon = \varepsilon_{out} + \varepsilon_{in}$ • round complexity: $k = k_{out} + k_{in}$ • randomness complexity: $r = r_{out} + r_{in}$
- proof length: $\ell = \ell_{out} + \underline{1} \cdot \ell_{in}$ (and similarly for prover time: $pt = pt_{out} + \underline{1} \cdot pt_{in}$)
- query complexity: $q = q_{in}$ (and similarly for verifier time: $vt = [\text{interaction of } V_{out}] + vt_{in}$)