Lecture 20

Foundations of Probabilistic Proofs
Fall 2020
Alessandro Chiesa

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? (polynomial proof length & constant query complexity)

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.
Intuitively, if we apply this to (i) and (ii) then we will get the best of both.
In particular this leads to a result known as the PCP Theorem:

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

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:

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

Let's start by building intuition on proof composition.

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]

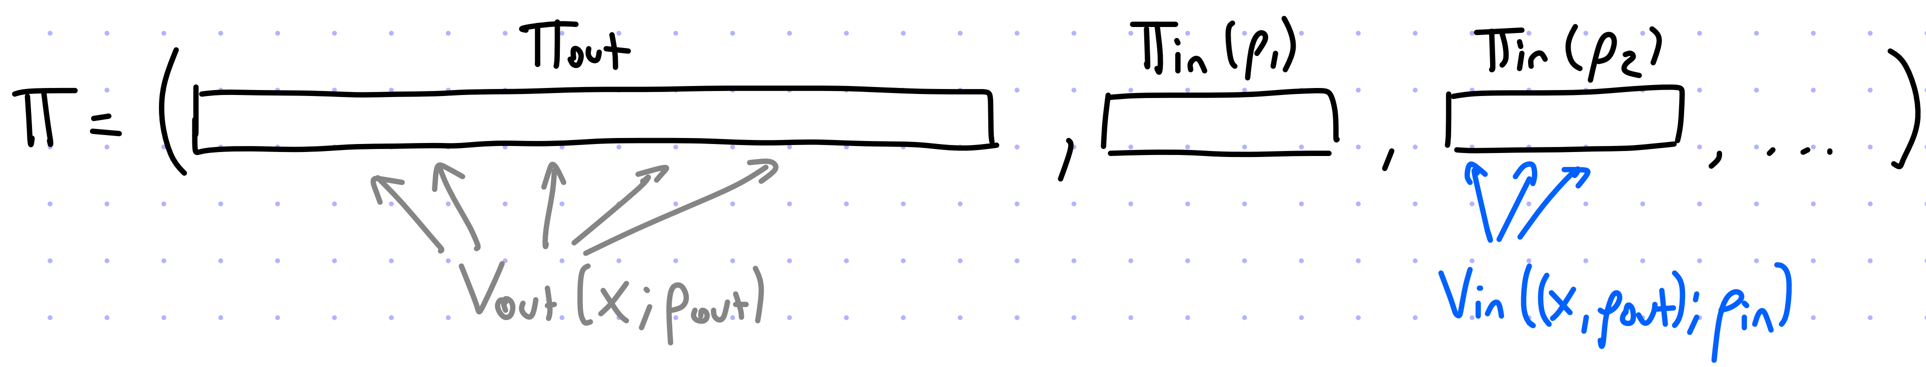

$P(x)$

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

$$\pi = \left(\underbrace{\pi_{out}}_{V_{out}(x;\rho_{out})}, \ \pi_{in}(\rho_1), \ \underbrace{\pi_{in}(\rho_2)}_{V_{in}((x,\rho_{out});\rho_{in})}, \ \dots \right)$$

possibly inconsistent choice across different ρ_{out}

- **Problem:** It could be that $\forall \rho_{out} \in \{0,1\}^{r_{out}}\ \exists \pi_{out}\ V_{out}^{\pi_{out}}(x;\rho_{out}) = 1$ (even when $x \notin L$).
  If so, the inner PCP is invoked on the true statement "$\exists \pi_{out}\ V_{out}^{\pi_{out}}(x;\rho_{out})$".

  Approach: Each inner PCP should be a "proof of proximity" for the corresponding local view.
  Compare: "is there a satisfying local view for (x,ρ_{out})?"
  vs. "is this local view (derived from the given π_{out}) satisfying for (x,ρ_{out})?"
  In particular each $\pi_{in}[\rho_{out}]$ is specifically about $\pi_{out}[Q_{out}(x,\rho_{out})]$: $V_{in}^{\underbrace{\pi_{out}[Q_{out}(x,\rho_{out})]}_{w_{in}},\pi_{in}[\rho_{out}]}(\underbrace{(x,\rho_{out})}_{x_{in}})$

- **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]

Soundness in a PCP states that the probability that a local view is accepting is small.

Robust soundness strengthens this to require that the probability that a local view is close to accepting is small. (In other words, whp a local view is far from accepting.)

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

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

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

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

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

① completeness: $\forall x \in L$, 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 robust soundness: $\forall x \notin L \ \forall \tilde{\pi} \ \Pr_{\rho}[\Delta(\tilde{\pi}[Q(x,\rho)], R(V)[(x,\rho)]) \leq \sigma] \leq \varepsilon_s$

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

(In fact also for any $\sigma \in [0, 1/q)$ because the strings being compared have q symbols.)

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).

Let $R = \{(x,w)| \dots\}$ 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 $R[x] = \emptyset$]

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

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

1. completeness: $\forall (x,w) \in R$, for $\pi := P(x,w)$, $\Pr_{\rho}[V^{w,\pi}(x;\rho)=1] \geq 1-\varepsilon_c$
2. proximity soundness: $\forall (x,w)$ if $\Delta(w,R[x]) \geq \delta$ then $\forall \tilde{\pi}$ $\Pr_{\rho}[V^{w,\tilde{\pi}}(x;\rho)=1] \leq \varepsilon_s$

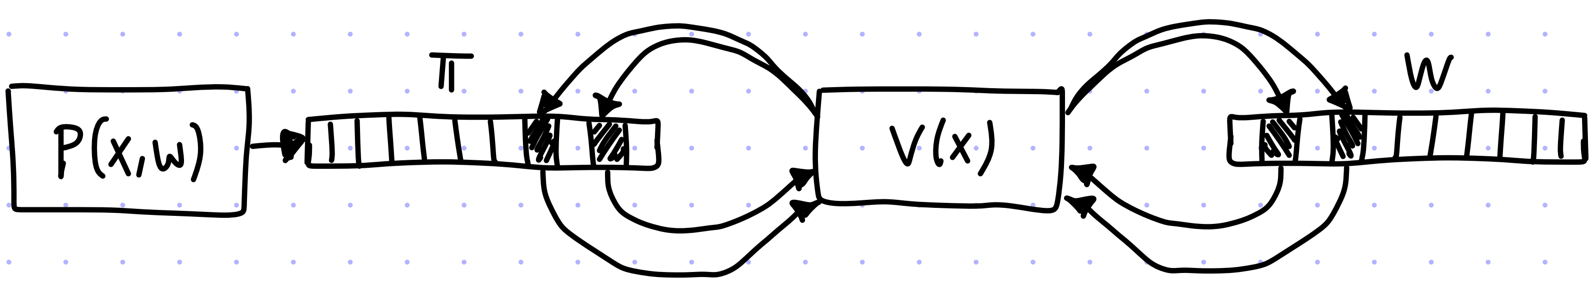

The query complexity counts queries to w & π.

Note: even if $x \in L(R)$, the PCPP verifier will still reject whp if w is far from $R[x]$

$\Rightarrow$ PCPPs are about proximity to valid witnesses not (just) about membership in $L(R)$.

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:

$$\Pi = \Big(\underbrace{\boxed{\qquad\qquad\qquad\qquad}}_{\Pi_{out}},\ \underbrace{\boxed{\qquad}}_{\Pi_{in}(\rho_1)},\ \underbrace{\boxed{\qquad}}_{\Pi_{in}(\rho_2)},\ \dots \Big)$$

V_{in} $((x, \rho_{out}); \rho_{in})$

$P(x)$

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

$V^{\Pi}(x)$

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

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}

<u>theorem</u>: 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: $vt = vt_{in}$)

Moreover:

rather than a language L

(a) if outer is a PCPP for a relation R with proximity δ_{out} then composed is a PCPP for R with proximity δ_{out}

<u>Why?</u> In construction and analysis consider (w, π_{out}) instead of just π_{out},

and for the soundness case consider w that is δ_{out}-far from $R[x]$ (rather than $x \notin L(R)$).

(b) if inner PCPP has robustness σ_{in} then the composed PCP has robustness σ_{in}

<u>Why?</u> Except with probability ε_{out}, the local view $\pi_{out}[Q_{out}(x, \rho_{out})]$ is σ_{out}-far from $R(V_{out})[(x, \rho_{out})]$.

If so, since $\sigma_{out} \geq \delta_{in}$, the local view $(\pi_{out}[Q_{out}(x, \rho_{out})]. \pi_{in}[\rho_{out}])[Q_{in}((x, \rho_{out}), \rho_{in})]$ is σ_{in}-close to accepting with probability at most ε_{in}.

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

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

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

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

  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.

And we can similarly compose IOPs via INTERACTIVE PROOF COMPOSITION, which is more efficient than its non-interactive counterpart thanks to interaction.

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:

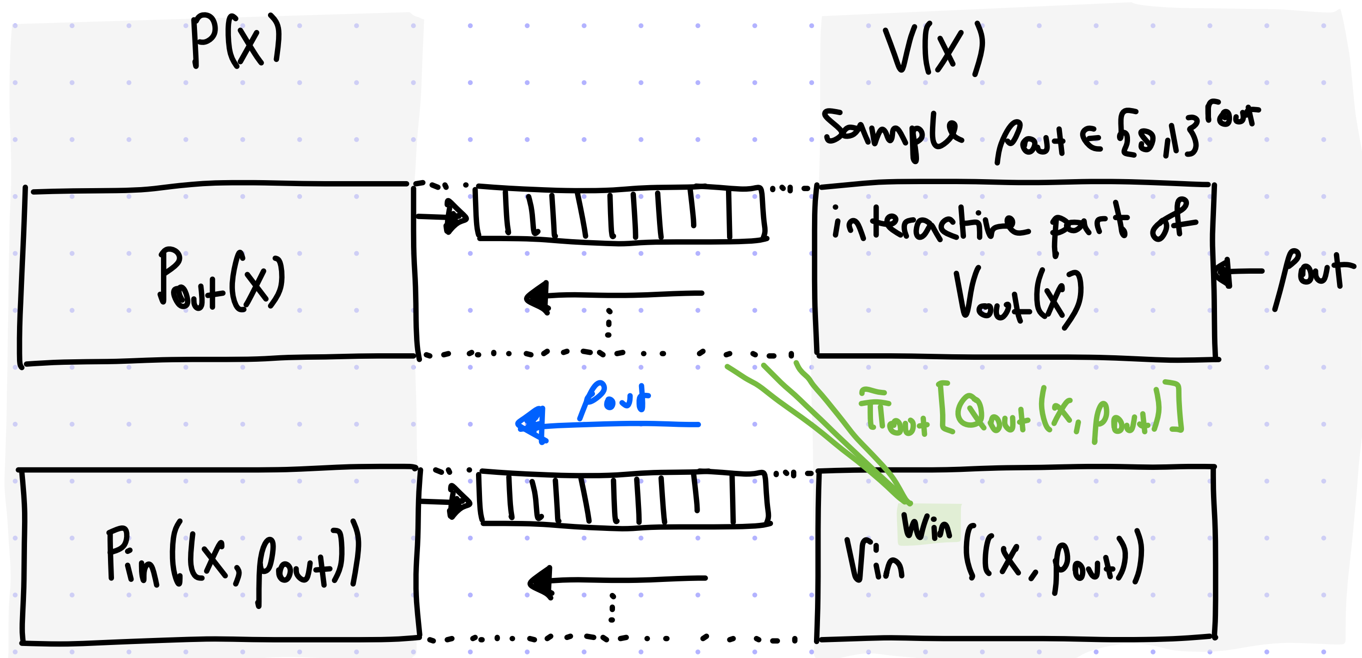

There is NO need to run inner IOP for every $\rho_{out} \in \{0,1\}^{r_{out}}$!

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: $vt = [\text{interaction of } V_{out}] + vt_{in}$)

Moreover: [similar statement as in the case of PCPs]