Foundations of Probabilistic Proofs

A course by **Alessandro Chiesa**

Lecture 19

Proof Composition & The PCP Theorem

Proof Composition

We saw techniques to achieve:

ⓘ polynomial proof length and polylogarithmic query complexity

ⓘⓘ exponential proof length and constant query complexity

How to achieve the best of both? (polynomial proof length and constant query complexity)

Proof Composition: 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 ⓘ and ⓘⓘ then we get the best of both.

This technique leads to

$$\underline{\text{PCP Theorem}}: \quad NP \subseteq PCP\begin{bmatrix} \varepsilon_c = 0 \\ \varepsilon_s = 1/2 \end{bmatrix}, \begin{matrix} \Sigma = \{0,1\} \\ \ell = \text{poly}(n) \end{matrix}, \begin{matrix} r = O(\log n) \\ q = O(1) \end{matrix}$$

Interactive Proof Composition: analogous technique that works for IOPs.

This technique leads to the optimal tradeoff between proof length and query complexity:

$$\underline{\text{theorem}}: \quad CSAT \in IOP\begin{bmatrix} \varepsilon_c = 0 \\ \varepsilon_s = 1/2 \end{bmatrix}, k = 3, \begin{matrix} \Sigma = \{0,1\} \\ \ell = O(n) \end{matrix}, \begin{matrix} r = O(\log n) \\ q = O(1) \end{matrix}$$

Today we study these techniques.

High-Level Plan

Ingredients: ⓘ outer PCP (P_{out}, V_{out}) for a relation R (with "good" proof length)

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

GOAL: PCP (P,V) for the relation R that inherits { outer's proof length ; inner's query complexity } .

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

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

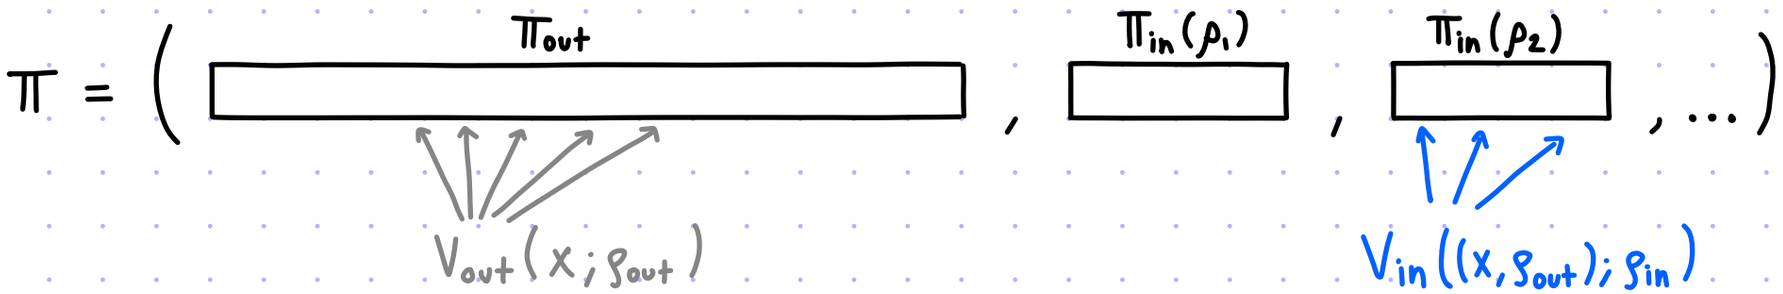

$P(x,w)$

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

   compute inner PCP for ρ_{out}

   $\pi_{in}(\rho_{out}) := P_{in}((x, \rho_{out}), \pi_{out}[Q_{out}(x, \rho_{out})])$.
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 $V_{in}^{\pi_{in}(\rho_{out})}(\overbrace{(x, \rho_{out})}^{x_{in}}) = 1$.

This plan has problems ...

Problems with the Plan

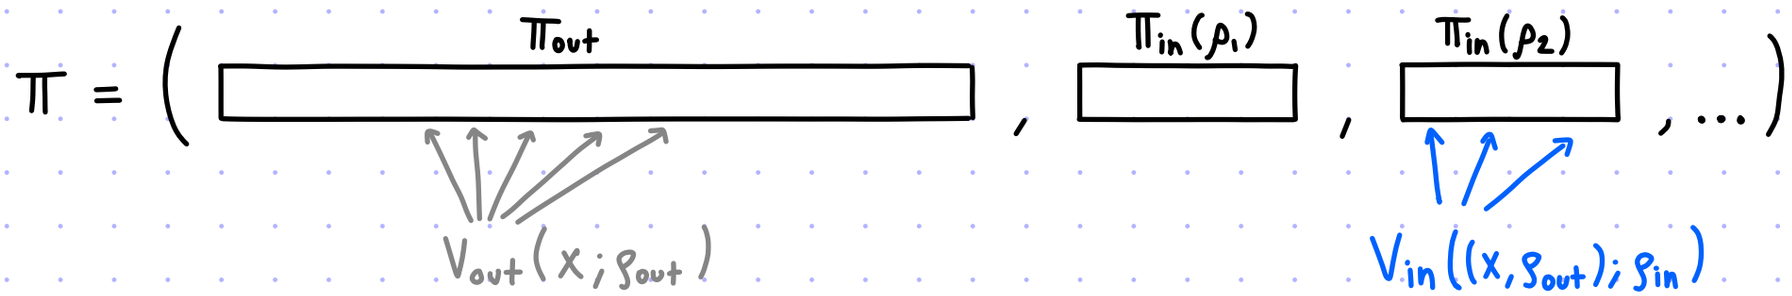

- **PROBLEM:** Even if $x \notin L$, it can be that $\forall \rho_{out} \in \{0,1\}^{r_{out}}$ $\exists$ ~~π_{out}~~ $V_{out}^{\pi_{out}}(x;\rho_{out}) = 1$. ← possibly inconsistent across different ρ_{out}

  If so, the inner PCP is invoked on the true statement "$\exists\, \pi_{out}\ V_{out}^{\pi_{out}}(x;\rho_{out}) = 1$".

Approach: Each inner PCP should be a "proof of proximity" for the corresponding local view.

Compare:
- "is there an accepting local view for (x,ρ_{out})?"
- "is THIS local view (derived from the given π_{out}) accepting for (x,ρ_{out})?"

Each $\pi_{in}(\rho_{out})$ will be 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 determine with few 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:

$x \notin L \rightarrow$ w.h.p. local view is far from ANY accepting local view

Robust PCPs [for outer PCP]

In a PCP, soundness states that $\Pr[\text{local view is accepting}]$ is small.

Robust soundness strengthens this: $\Pr[\text{local view is close to accepting}]$ is small.

In other words, whp a local view is far from accepting.

We restrict attention to **non-adaptive verifiers**:

$V^{\pi}(x;\rho) = D(S(x,\rho), \pi[Q(x,\rho)])$ where S, Q, D are the state, query, decision algorithms of V.

The relation of accepting local views for $V=(S,Q,D)$ is:

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

Given an instance s, $R(V)[s] := \{a \mid (s,a) \in R(V)\}$.

def: (P,V) is a PCP system for a relation R with **robustness parameter σ** if:

① **completeness:** $\forall (x,w) \in R \quad \Pr\left[V^{\pi}(x)=1 \;\middle|\; \pi \leftarrow P(x,w)\right] \geq 1-\varepsilon_c$.

② **robust soundness:** $\forall x \notin L(R) \quad \forall \tilde{\pi} \quad \Pr_{\rho}\left[\Delta(\tilde{\pi}[Q(x,\rho)], R(V)[S(x,\rho)]) \leq \sigma\right] \leq \varepsilon_s$.

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

(Also for every $\sigma \in [0, 1/q)$ because a local view has q query symbols.)

PCPs of Proximity [for inner PCP]

In a **PCP of proximity (PCPP)** for a relation R the verifier receives:

- an instance x
- query access to a candidate witness w
- query access to a PCP string π

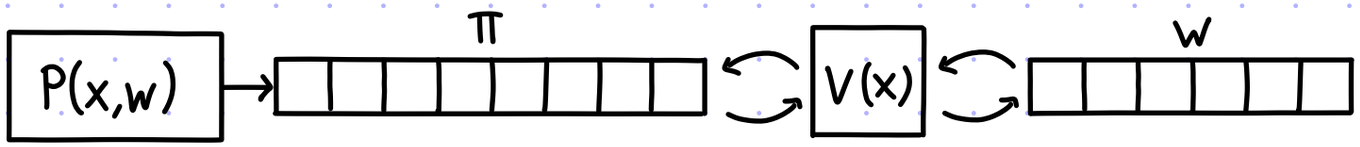

GOAL: convince the verifier that w is close to some valid witness in $R[x] := \{w \mid (x,w) \in R\}$.

(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:

① completeness: $\forall (x,w) \in R \quad \Pr\left[V^{w,\pi}(x)=1 \;\middle|\; \pi \leftarrow P(x,w)\right] \geq 1-\varepsilon_c$.

② proximity soundness: $\forall (x,w)$ if $\Delta(w,R[x]) \geq \delta$ then $\forall \tilde{P} \quad \Pr\left[V^{w,\tilde{\pi}}(x)=1 \;\middle|\; \tilde{\pi} \leftarrow \tilde{P}\right] \leq \varepsilon_s$.

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

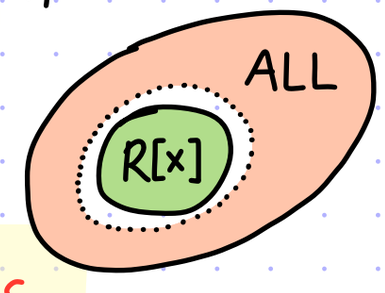

Equivalently:
$\Delta(w,R[x]) \geq \delta \rightarrow \forall \tilde{\pi} \; \Pr[V^{w,\tilde{\pi}}(x)=1] \leq \varepsilon_s$

EFFICIENCY: proof length measures $|\pi|$ (over a given alphabet)

but query complexity counts queries to w and π.

NOTE: if $x \in L(R)$ then the PCPP verifier rejects w.h.p. if w is far from R[x]

⇒ 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:

$\Pi = ($ 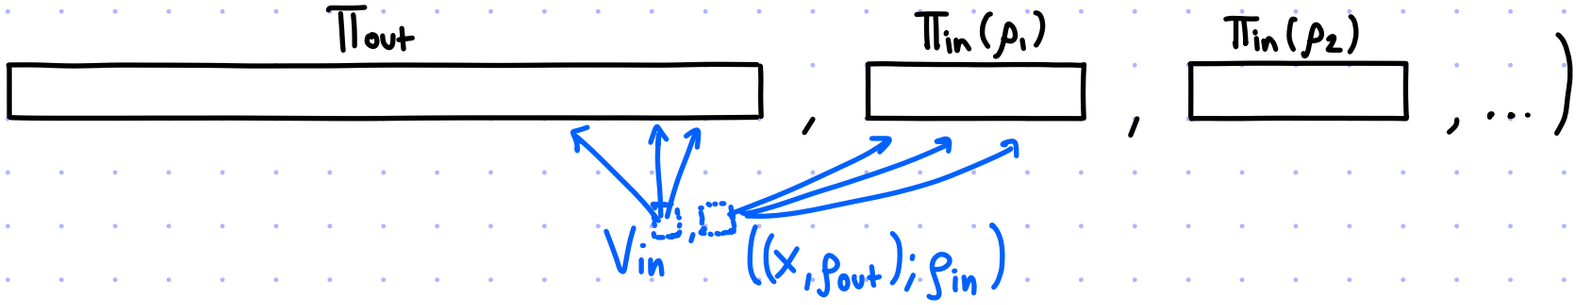 $)$

$P(x)$

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

   $\pi_{in}[\rho_{out}] := P_{in}(S_{out}(x, \rho_{out}), \pi_{out}[Q_{out}(x, \rho_{out})])$
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 $V_{in}^{\overbrace{\pi_{out}[Q_{out}(x,\rho_{out})]}^{w_{in}}, \pi_{in}[\rho_{out}]}(\underbrace{S_{out}(x, \rho_{out})}_{x_{in}}) = 1$.

<u>claim</u>: The soundness error is $\varepsilon_{out} + \varepsilon_{in}$.

If $x \notin L$ then, except w.p. ε_{out} over ρ_{out}, the local view $\pi_{out}[Q_{out}(x, \rho_{out})]$ is σ_{out}-far from $R(V_{out})[S_{out}(x, \rho_{out})]$.

If so (and $\sigma_{out} \geq \delta_{in}$) then V_{in} accepts w.p. ε_{in} over ρ_{in}.

Proof Composition Theorem

$P(x)$

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

   $\pi_{in}[\rho_{out}] := P_{in}(S_{out}(x, \rho_{out}), \pi_{out}[Q_{out}(x, \rho_{out})])$
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 $V_{in}^{\underbrace{\pi_{out}[Q_{out}(x,\rho_{out})]}_{w_{in}}, \pi_{in}[\rho_{out}]}(\underbrace{S_{out}(x,\rho_{out})}_{x_{in}}) = 1$.

theorem: Consider these 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}

Then we obtain a PCP (P,V) for the language L with:

- soundness error: $\sigma_{out}(x) \geq \delta_{in}(x_{in}) \rightarrow \varepsilon(x) = \varepsilon_{out}(x) + \varepsilon_{in}(x_{in})$
- proof length: $\ell(x) = \ell_{out}(x) + 2^{r_{out}(x)} \cdot \ell_{in}(x_{in})$
- query complexity: $q(x) = q_{in}(x_{in})$
- randomness complexity: $r(x) = r_{out}(x) + r_{in}(x_{in})$
- prover time: $pt(x) = pt_{out}(x) + 2^{r_{out}(x)} \cdot (st_{out}(x) + qt_{out}(x) + pt_{in}(x_{in}))$
- verifier time: $vt(x) = st_{out}(x) + qt_{out}(x) + vt_{in}(x_{in})$

Variations on Proof Composition

lemma: if (P_{in}, V_{in}) has robustness σ_{in} then (P,V) has robustness σ_{in}

proof:

If $x \notin L$ then, except w.p. ε_{out} over ρ_{out}, the local view $\pi_{out}[Q_{out}(x, \rho_{out})]$ is σ_{out}-far from $R(V_{out})[S_{out}(x, \rho_{out})]$.

If so (and $\sigma_{out} \geq \delta_{in}$) then the local view

$$\left(\pi_{out}[Q_{out}(x,\rho_{out})], \pi_{in}[\rho_{out}]\right)[Q_{in}(S_{out}(x,\rho_{out}), \rho_{in})]$$

is σ_{in}-far from $R(V_{in})[S_{in}(S_{out}(x,\rho_{out}), \rho_{in})]$ except w.p. ε_{in} over ρ_{in}. ■

lemma: if (P_{out}, V_{out}) is a PCPP for a relation R with proximity δ_{out} then (P,V) is a PCPP for R with proximity δ_{out}

proof: In the construction and analysis consider local views of (w, π_{out}) rather than of π_{out}.

$P(x,w)$

1. Compute outer PCP: $\pi_{out} := P_{out}(x,w)$
2. For each $\rho_{out} \in \{0,1\}^{r_{out}}$:
   compute inner PCPP for ρ_{out} as
   $\pi_{in}[\rho_{out}] := P_{in}(S_{out}(x,\rho_{out}), (w, \pi_{out})[Q_{out}(x,\rho_{out})])$
3. Output $\pi := (\pi_{out}, (\pi_{in}[\rho_{out}])_{\rho_{out} \in \{0,1\}^{r_{out}}})$.

$V^{w,\pi}(x)$

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

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

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

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

(accepting ↑ local views for $S(x,\rho)$)

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

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

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

Example: If $R = \{((\mathbb{F},L,d),f) \mid f \in RS[\mathbb{F},L,d]\}$ then we get an IOPP for the Reed-Solomon code.
FRI is an example.

We can similarly compose IOPs via INTERACTIVE PROOF COMPOSITION.
It is more efficient than its non-interactive counterpart thanks to interaction.

Interactive Proof Composition [1/2]

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}

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

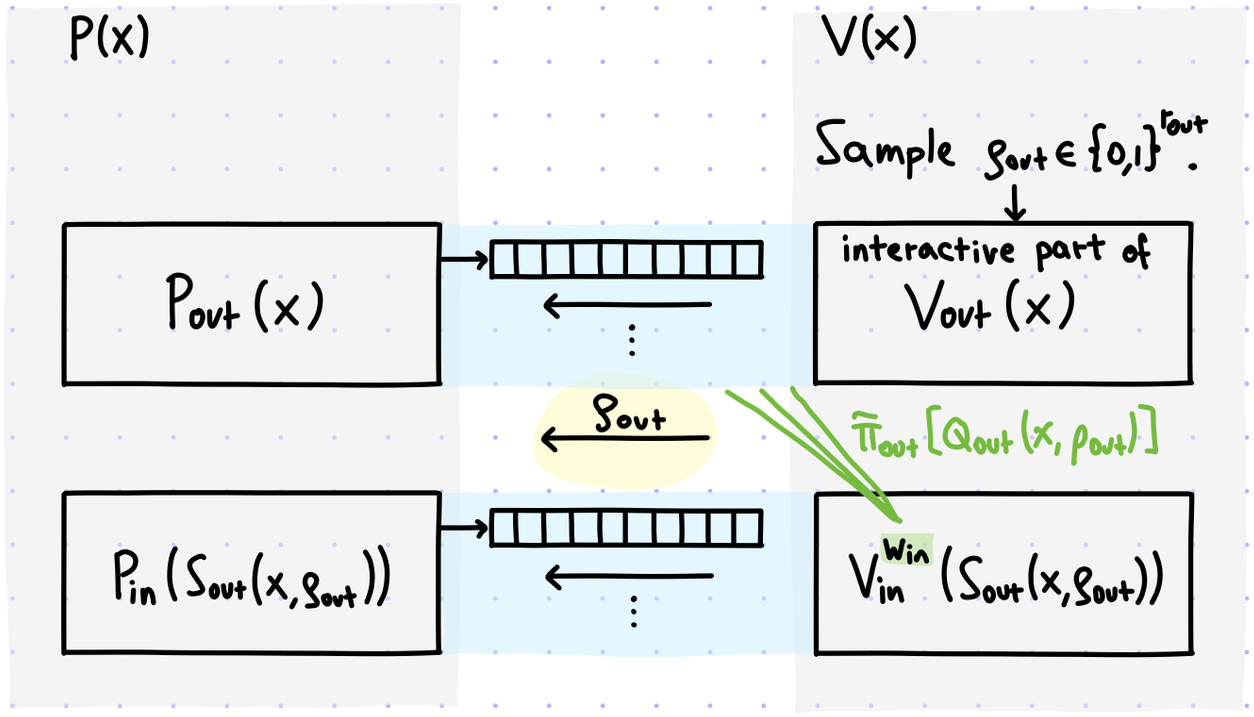

There is no need to run the inner IOP for every $\rho_{out} \in \{0,1\}^{r_{out}}$:

the IOP verifier tells the IOP prover which ρ_{out} it sampled.

Interactive Proof Composition [2/2]

theorem: Consider these 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}

Then we obtain an IOP (P,V) for the language L with:

- soundness error: $\sigma_{out}(x) \geq \delta_{in}(x_{in}) \rightarrow \varepsilon(x) = \varepsilon_{out}(x) + \varepsilon_{in}(x_{in})$
- round complexity: $k(x) = k_{out}(x) + k_{in}(x_{in})$
- proof length: $\ell(x) = \ell_{out}(x) + 1 \cdot \ell_{in}(x_{in})$
- query complexity: $q(x) = q_{in}(x_{in})$
- randomness complexity: $r(x) = r_{out}(x) + r_{in}(x_{in})$
- prover time: $pt(x) = pt_{out}(x) + 1 \cdot (st_{out}(x) + qt_{out}(x) + pt_{in}(x_{in}))$
- verifier time: $vt(x) = st_{out}(x) + qt_{out}(x) + vt_{in}(x_{in})$

lemma: if (P_{in}, V_{in}) has robustness σ_{in} then (P,V) has robustness σ_{in}

lemma: if (P_{out}, V_{out}) is an IOPP for a relation R with proximity δ_{out} then (P,V) is an IOPP for R with proximity δ_{out}

PCP Theorem via Proof Composition [1/3]

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

Below we implicitly require $\varepsilon_c = 0, \varepsilon_s = 1/2, \Sigma = \{0,1\}$ (we omit them to reduce clutter).

PROOF ATTEMPT Apply (non-interactive) proof composition with:

- outer PCP: robust variant of PCP for NP with proof length poly(n) and query complexity poly(log n)

  $CSAT \in PCP[\ell_{out} = \text{poly}(n), q_{out} = \text{poly}(\log n), r_{out} = O(\log n), s_{out} = \text{poly}(\log n), \sigma_{out} = \Omega(1)]$

- inner PCPP: proximity variant of the PCP for NP with proof length exp(n) and query complexity O(1)

  $R(V_{out}) \in PCPP[\ell_{in} = \exp(n_{in}), q_{in} = O(1), r_{in} = \text{poly}(n_{in}), \delta_{in} = O(1)]$

We ensure that $\sigma_{out} \geq \delta_{in}$ and set $n_{in} := s_{out}(n)$. Proof composition yields a PCP for CSAT with:

$$CSAT \in PCP \begin{bmatrix} \ell = \ell_{out} + 2^{r_{out}} \cdot \ell_{in} = \text{poly}(n) + 2^{O(\log n)} \cdot \exp(\text{poly}(\log n)) = n^{\text{poly}(\log n)} \\ q = q_{in} = O(1) \\ r = r_{out} + r_{in} = O(\log n) + \text{poly}(\text{poly}(\log n)) = \text{poly}(\log n) \end{bmatrix} \leftarrow \text{TOO LONG!}$$

IDEA Step 1: compose the outer PCP with itself to obtain a smaller state size

Step 2: compose the resulting PCP with the inner PCPP

This requires starting from a robust PCPP.

PCP Theorem via Proof Composition [2/3]

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

PART 1 OF PROOF Apply (non-interactive) proof composition with:

- **outer PCP:** robust variant of PCP for NP with proof length poly(n) and query complexity poly(log n)

like in prior slide: $CSAT \in PCP[\ell_{out} = \text{poly}(n), q_{out} = \text{poly}(\log n), r_{out} = O(\log n), s_{out} = \text{poly}(\log n), \sigma_{out} = \Omega(1)]$

- **inner PCPP:** robust & proximity variant of the PCP for NP with proof length poly(n) and query complexity poly(log n)

$R(V_{out}) \in PCPP[\ell_{in} = \text{poly}(n_{in}), q_{in} = \text{poly}(\log n_{in}), r_{in} = O(\log n_{in}), s_{in} = \text{poly}(\log n_{in}), \delta_{in} = O(1), \sigma_{in} = \Omega(1)]$

We ensure that $\sigma_{out} \geq \delta_{in}$ and set $n_{in} := s_{out}(n)$. Proof composition yields a PCP for CSAT with:

$$CSAT \in PCP\left[\begin{array}{l} \ell = \ell_{out} + 2^{r_{out}} \cdot \ell_{in} = \text{poly}(n) + 2^{O(\log n)} \cdot \text{poly}(\text{poly}(\log n)) = \text{poly}(n) \\ q = q_{in} = \text{poly}(\log \text{poly}(\log n)) = \text{poly}(\log\log n) \\ r = r_{out} + r_{in} = O(\log n) + O(\log \text{poly}(\log n)) = O(\log n) \\ s = s_{in} = \text{poly}(\log \text{poly}(\log n)) = \text{poly}(\log\log n) \\ \sigma = \sigma_{in} = \Omega(1) \end{array}\right]$$

The composed PCP will act as the outer PCP in the next composition.

We used the fact that if the inner PCPP is robust then so is the composed PCP.

We keep track of the state size for the composed PCP (it is $s_{in}(s_{out}(n))$).

PCP Theorem via Proof Composition [3/3]

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

Part 2 of Proof Apply (non-interactive) proof composition with:

- **outer PCP:** robust PCP for NP obtained from the first composition

  $CSAT \in PCP[\ell_{out} = poly(n), q_{out} = poly(\log\log n), r_{out} = O(\log n), s_{out} = poly(\log\log n), \sigma_{out} = \Omega(1)]$

- **inner PCPP:** proximity variant of the PCP for NP with proof length exp(n) and query complexity O(1)

  $R(V_{out}) \in PCPP[\ell_{in} = \exp(n_{in}), q_{in} = O(1), r_{in} = poly(n_{in}), \delta_{in} = O(1)]$

We ensure that $\sigma_{out} \geq \delta_{in}$ and set $n_{in} := s_{out}(n)$. Proof composition yields a PCP for CSAT with:

$$CSAT \in PCP\left[\begin{array}{l} \ell = \ell_{out} + 2^{r_{out}} \cdot \ell_{in} = poly(n) + 2^{O(\log n)} \cdot \exp(poly(\log\log n)) = poly(n) \\ q = q_{in} = O(1) \\ r = r_{out} + r_{in} = O(\log n) + poly(poly(\log\log n)) = O(\log n) \end{array}\right]$$

Remarks on Proof Composition

Summary:

Step 1

robust PCP	∘	robust PCPP	→	robust PCP	∘	PCPP	→	PCP
$\ell = poly(n)$		$\ell = poly(n)$		$\ell = poly(n)$		$\ell = exp(n)$		$\ell = poly(n)$
$q = poly(\log n)$		$q = poly(\log n)$		$q = poly(\log\log n)$		$q = O(1)$		$q = O(1)$

Step 2

Why not compose the robust PCPP with itself 3 (or more) times?

Step 1

robust PCP	∘	robust PCPP	→	robust PCP	∘	robust PCPP	→	PCP
$\ell = poly(n)$		$\ell = poly(n)$		$\ell = poly(n)$		$\ell = poly(n)$		$\ell = poly(n)$
$q = poly(\log n)$		$q = poly(\log n)$		$q = poly(\log\log n)$		$q = poly(\log n)$		$q = poly(\log\log\log n)$

Step 2

query complexity decreases but not to constant (& ditto for more compositions)

Proximity variant of the PCP Theorem:

theorem: $\forall \delta > 0 \quad NP \subseteq PCPP[\varepsilon_c = 0, \varepsilon_s = \frac{1}{2}, \Sigma = \{0,1\}, \ell = poly(n), q = O(1), r = O(\log n), \delta]$

proof: Perform the 2-step composition used to prove the PCP Theorem, with a modification. In the first composition, set the outer PCP to be a robust & proximity variant of the PCP for NP with proof length poly(n) and query complexity poly(log n). (Eg the same as the inner PCPP.) Both compositions preserve the proximity parameter. ■

Robust variant of the PCP Theorem: straightforward because $q = O(1)$ (e.g. use an error-correcting code).

Bibliography

PCP Theorem

- New short cut found for long math proofs, New York Times 1992.
- PCP theorem, Wikipedia.
- [ALMSS 1998]: Proof verification and the hardness of approximation problems, by Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, Mario Szegedy.

PCP composition

- [BGHSV 2005]: Robust PCPs of proximity, shorter PCPs and applications to coding, by Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, Salil Vadan.
- [DR 2006]: Assignment testers: towards a combinatorial proof of the PCP theorem, by Irit Dinur, Omer Reingold.

IOP composition

- [BCGRS 2016]: Interactive oracle proofs with constant rate and query complexity, by Eli Ben-Sasson, Alessandro Chiesa, Ariel Gabizon, Michael Riabzev, Nick Spooner.
- [RR 2020]: Local proofs approaching the witness length, by Noga Ron-Zewi, Ron Rothblum. (▸Video).