Lecture 22

Foundations of Probabilistic Proofs
Fall 2020
Alessandro Chiesa

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).

For a relation $R = \{(x,w) | \cdots\}$, the language of R is $L(R) = \{x | \exists w \text{ s.t. } (x,w) \in R\}$ and the valid witnesses of an instance x is $R[x] = \{w | (x,w) \in R\}$. [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 $\forall \tilde{\pi}$ $\Pr_\rho[V^{w,\tilde{\pi}}(x;\rho)=1] \leq \varepsilon_s$

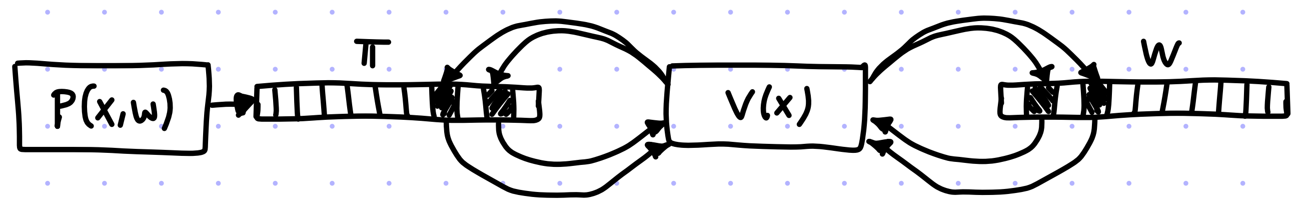

We show how to construct two types of PCPPs

theorem: $\forall \delta > 0$ $\text{QESAT}(\mathbb{F}_2) \in \text{PCPP}[\varepsilon_c = 0, \varepsilon_s = 1/2, \Sigma = \{0,1\}, \ell = \exp(n), q = O(\frac{1}{\delta}), r = \text{poly}(n), \delta]$

theorem: $\forall \delta > 0$ $\text{QESAT}(\mathbb{F}_2) \in \text{PCPP}[\varepsilon_c = 0, \varepsilon_s = 1/2, \Sigma = \{0,1\}, \ell = \text{poly}(n), q = O\left(\frac{\text{polylog}(n)}{\delta}\right), r = O(\log n), \delta]$

Here $\text{QESAT}(\mathbb{F})$ is the relation $\{((p_1,\ldots,p_m),a) \mid p_1,\ldots,p_m \in \mathbb{F}^{\leq 2}[X_1,\ldots,X_n], a:[n] \to \mathbb{F}, p_1(a) = \cdots = p_m(a) = 0\}$

Easy: from PCPP to PCP

lemma: Fix any proximity parameter δ. Suppose that a relation R in $\mathrm{PCPP}[\varepsilon_c, \varepsilon_s, \Sigma, \ell, r, q, \delta]$. Then the language $L(R)$ is in $\mathrm{PCP}[\varepsilon_c, \varepsilon_s, \Sigma, \ell' = |w| + \ell, r, q]$.

proof: Let (P_{px}, V_{px}) be the PCPP for R. We construct the PCP (P, V) for $L(R)$ as follows:

$P(x)$

1. Find witness w for x. (Or receive it as input.)
2. Compute proximity proof: $\pi_{px} := P_{px}(x, w)$
3. Output $\pi := (w, \pi_{px})$.

$\pi = \left(\boxed{\quad w \quad}, \boxed{\qquad\quad \pi_{px} \qquad\quad} \right)$

$V^{\pi}(x) =$ Check that $V_{px}^{w, \pi_{px}}(x)$ accepts.

Completeness: If $x \in L(R)$ then $\exists\, w$ s.t. $(x, w) \in R$ so, for $\pi_{px} := P_{px}(x, w)$, $V_{px}^{w, \pi_{px}}(x) = 1$ w.p. $\geq 1 - \varepsilon_c$.

Soundness: If $x \notin L(R)$ then all candidate witnesses are far from $R[x]$:

$$\forall \tilde{w} \quad \Delta(\tilde{w}, R[x]) = 1 \geq \delta \text{ so } \forall \tilde{\pi} = (\tilde{w}, \tilde{\pi}_{px}) \quad V_{px}^{\tilde{w}, \tilde{\pi}_{px}}(x) = 1 \text{ w.p.} \leq \varepsilon_s. \qquad \blacksquare$$

Thus PCPPs are "stronger" than PCPs (even though PCPPs are about proximity rather than satisfiability). The extra power is when $x \in L(R)$: the PCPP verifier rejects whp if the given w is far from $R[x]$.

So we have to work at least as hard to construct PCPPs as we did for PCPs.

Harder: from PCP to PCPP

We recycle: we modify PCPs for languages $L(R)$ into PCPPs for the corresponding relation R.

Recall that our recipe to construct PCPs so far has been to set $\pi = (\pi_a, \pi_{sat})$ where

① π_a is (allegedly) the encoding of a candidate witness [belongs to $S := \{Enc(z)\}_z$]

② if π_a is close to $Enc(a)$ for some a, π_{sat} facilitates checking that a is a valid witness.

In fact, the analysis showed that if the PCP verifier $V_{PCP}^{(\pi_a,\pi_{sat})}(x)=1$ accepts whp then not only we learn that $x \in L(R)$ but also that π_a is close to $Enc(\tilde{w})$ for $(x,\tilde{w}) \in R$.

def: V_{PCP} is $(\varepsilon_{PCP}, \delta_{PCP}, Enc)$-sound if $\forall\ (\tilde{\pi}_a, \tilde{\pi}_{sat})$ $\Pr[V_{PCP}^{(\tilde{\pi}_a,\tilde{\pi}_{sat})}(x)=1] > \varepsilon_{PCP}$ implies $\exists\, \tilde{w} \in R[x]$ s.t. $\Delta(\tilde{\pi}_a, Enc(\tilde{w})) < \delta$.

This leads to a template construction of a PCPP (P,V) for R from a PCP (P_{PCP}, V_{PCP}) for $L(R)$ as above:

$P(x,w)$

1. Compute the PCP (π_a, π_{sat}) output by $P_{PCP}(x)$ when using the witness w.
2. Compute proof for encoding consistency test: $\pi_e := P_e(w, \pi_a)$.

$V^{w,\, \pi_{px} = (\pi_a, \pi_{sat}, \pi_e)}(x)$

1. Check that $V_{PCP}^{(\pi_a,\pi_{sat})}(x)=1$
2. Check that $V_e^{w,\pi_a,\pi_e} = 1$

The encoding consistency test satisfies the following property:

if w is δ-far from $\tilde{w}$ and π_a is δ_{PCP}-close to $Enc(\tilde{w})$ then $\forall \pi_e$ $V_e^{w,\pi_a,\pi_e}(x)=1$ w.p. $\leq \varepsilon_e(\delta, \delta_{PCP})$.

Consistency Test via Local Decoders

We say that D is a local decoder for Enc with decoding radius δ_{LD} and error probability ε_{LD} if

① $\forall a \ \forall i \in [n] \ \Pr[D^{Enc(a)}(i) = a_i] = 1$

② if $\tilde{\pi}$ is δ_{LD}-close to $Enc(a)$ then $\forall i \in [n] \ \Pr[D^{\tilde{\pi}}(i) \neq a_i] \leq \varepsilon_{LD}$

We can use local decoders to do an encoding consistency test without an auxiliary proof π_e:

<u>lemma</u>: Suppose that $L(R) \in PCP[(\varepsilon_{PCP}, \delta_{PCP}, Enc), \Sigma, \ell, q, r]$ and Enc has a local decoder with decoding radius $\delta_{LD} \geq \delta_{PCP}$ and error probability ε_{LD}. Then $\forall \delta, \varepsilon > 0$

$$R \in PCPP\left[\varepsilon' \leq \max\{\varepsilon_{PCP}, \varepsilon\}, \Sigma, \ell, q' = q + O\left(\frac{\log 1/\varepsilon}{(1-\varepsilon_{LD})\delta}\right) q_{LD}, r' = r + O\left(\frac{\log 1/\varepsilon}{(1-\varepsilon_{LD})\delta}\right)(\log|w| + r_{LD}), \delta\right]$$

Here is the construction of the PCPP where we set $t = O\left(\frac{\log 1/\varepsilon}{(1-\varepsilon_{LD})\delta}\right)$:

$P(x, w)$

1. Compute the PCP (π_a, π_{sat}) output by $P_{PCP}(x)$ when using the witness w.
2. Output $\pi_{px} := (\pi_a, \pi_{sat}, \bot)$.

$V^{w, \pi_{px} = (\pi_a, \pi_{sat}, \bot)}(x)$

1. Check that $V_{PCP}^{(\pi_a, \pi_{sat})}(x) = 1$
2. Sample $i_1, \ldots, i_t \in [|w|]$ and check that $\forall j \in [t] \ D^{\pi_a}(i_j) = w_{i_j}$

Consistency Test via Local Decoders

$P(x,w)$

1. Compute the PCP (π_a, π_{sat}) output by $P_{PCP}(x)$ when using the witness w.
2. Output $\pi_{px} := (\pi_a, \pi_{sat}, \perp)$.

$V^{w, \pi_{px} = (\pi_a, \pi_{sat}, \perp)}(x)$

1. Check that $V_{PCP}^{(\pi_a, \pi_{sat})}(x) = 1$
2. Sample $i_1, \dots, i_t \in [|w|]$ and check that $\forall j \in [t] \ D^{\pi_a}(i_j) = w_{i_j}$

Analysis

Completeness: if $(x,w) \in R$ then (i) by completeness of (P_{PCP}, V_{PCP}), $V_{PCP}^{(\pi_a, \pi_{sat})}(x) = 1$ w.p. 1 and (ii) since $\pi_a = Enc(w)$, by completeness of D, $\forall i \in [|w|] \ \Pr[D^{\pi_a}(i) = w_i] = 1$.

Soundness: if w is δ-far from $R[x]$ then $\forall (\tilde{\pi}_a, \tilde{\pi}_{sat})$ either

(i) $V_{PCP}^{(\tilde{\pi}_a, \tilde{\pi}_{sat})}(x) = 1$ w.p. $\leq \varepsilon_{PCP}$

OR

(ii) $x \in L(R)$ and $\tilde{\pi}_a$ is δ_{PCP}-close to $Enc(\tilde{w})$ for some $\tilde{w} \in R[x]$. [As V_{PCP} is $(\varepsilon_{PCP}, \delta_{PCP}, Enc)$-sound.]

In the latter case, since $\delta_{PCP} \leq \delta_{LD}$, $\forall i \in [|w|] \ \Pr[D^{\tilde{\pi}_a}(i) = \tilde{w}_i] \geq 1 - \varepsilon_{LD}$.

Since w is δ-far from $R[x]$, w is also δ-far from $\tilde{w} \in R[x]$, i.e., $\Pr_i[w_i \neq \tilde{w}_i] \geq \delta$.

We deduce that $V^{w, (\tilde{\pi}_a, \tilde{\pi}_{sat}, \perp)}(x) = 1$ w.p. $\leq \max\{\varepsilon_{PCP}, (1 - (1-\varepsilon_{LD})\delta)^t\}$.

So $t = O\left(\frac{\log 1/\varepsilon}{(1-\varepsilon_{LD})\delta}\right)$ gives $\max\{\varepsilon_{PCP}, \varepsilon\}$.

Exponential-Size Constant-Query PCPP

We constructed an exp-size constant-query PCP for QESAT($\mathbb{F}$) (quadratic equations over $\mathbb{F}$).

The encoding that we used was linear extensions: $\mathrm{Enc}: \mathbb{F}^n \to \mathbb{F}$ where $\mathrm{Enc}(a) = \{\langle a, c\rangle\}_{c \in \mathbb{F}^n}$.

Observe that:

- the soundness analysis showed that the PCP is $(\varepsilon_{PCP} = O(1), \delta_{PCP}, \mathrm{Enc})$-sound $\forall\, \delta_{PCP} \leq \frac{1}{2} \cdot \left(1 - \frac{1}{|\mathbb{F}|}\right)$
- Enc has a local decoder (in fact, a local corrector):

$$D^{\tilde{\pi}}(i) := \text{Sample } r_1, \ldots, r_t \in \mathbb{F}^n \text{ and return } \mathrm{plurality}_{j \in [t]} \{\tilde{\pi}(e_i + r_j) - \tilde{\pi}(r_j)\}$$

If $\tilde{\pi}$ is δ_{LD}-close to $\mathrm{Enc}(a)$ then $\Pr[D^{\tilde{\pi}}(i) \neq a_i] \leq \exp(-(1 - 2\delta_{LD}) \cdot t) \leq \varepsilon_{LD}$ for $t = \Theta\left(\frac{\log \frac{1}{\varepsilon_{LD}}}{1 - 2\delta_{LD}}\right)$.

Hence, focusing for simplicity on $\mathbb{F} = \mathbb{F}_2$, we can apply the lemma to this PCP and this local decoder:

theorem: $\forall\, \delta > 0$ $\mathrm{QESAT}(\mathbb{F}_2) \in \mathrm{PCPP}[\varepsilon_c = 0, \varepsilon_s = 1/2, \Sigma = \{0,1\}, \ell = \exp(n), q = O(\frac{1}{\delta}), r = \mathrm{poly}(n), \delta]$

Polynomial-Size Polylog-Query PCPP

We constructed an poly-size polylog-query PCP for QESAT($\mathbb{F}$) (quadratic equations over $\mathbb{F}$).

The encoding that we used was multivariate low-degree extensions:

$\mathrm{Enc}(a): \mathbb{F}^{\frac{\log n}{\log|H|}} \to \mathbb{F}$ where $\mathrm{Enc}(a) :=$ "$(\mathbb{F}, H, \frac{\log n}{\log|H|})$-extension of a" [which has total degree $d := \frac{\log n}{\log|H|} \cdot |H|$]

Observe that:

- the soundness analysis showed that the PCP is $(\varepsilon_{PCP} = O(1), \delta_{PCP}, \mathrm{Enc})$-sound $\forall\, \delta_{PCP} \leq \frac{1}{2} \cdot \left(1 - \frac{d}{|\mathbb{F}|}\right)$
- Enc has a local decoder (in fact, a local corrector):

magical constants about derivatives

$$D^{\tilde{\pi}}(i) := \text{Sample } r_1, \dots, r_t \in \mathbb{F}^{\frac{\log n}{\log|H|}} \text{ and return plurality}_{j \in [t]} \left\{ \sum_{k=1}^{d+1} c_i\, \tilde{\pi}(e_i + k \cdot r_j) \right\}$$

If $\tilde{\pi}$ is δ_{LD}-close to $\mathrm{Enc}(a)$ then $\Pr[D^{\tilde{\pi}}(i) \neq a_i] \leq \exp(-(1-(d+1)\delta_{LD}) \cdot t) \leq \varepsilon_{LD}$ for $t = \Theta\left(\frac{\log \frac{1}{\varepsilon_{LD}}}{1-(d+1)\delta_{LD}}\right)$.

Hence, focusing for simplicity on $\mathbb{F} = \mathbb{F}_2$, we can apply the lemma to this PCP and this local decoder:

theorem: $\forall\, \delta > 0$ QESAT($\mathbb{F}_2$) $\in$ PCPP$\left[\varepsilon_c = 0, \varepsilon_s = \frac{1}{2}, \Sigma = \{0,1\}, \ell = \mathrm{poly}(n), q = O\left(\frac{\mathrm{polylog}(n)}{\delta}\right), r = O(\log n), \delta\right]$

Robustness and Proximity

In fact we will need a PCPP that is also robust:

theorem: $\forall \delta > 0$ QESAT$(\mathbb{F}_2) \in$ PCPP$\left[\varepsilon_c = 0, \varepsilon_s = 1/2, \Sigma = \{0,1\}, \ell = \text{poly}(n), q = O\left(\frac{\text{polylog}(n)}{\delta}\right), r = O(\log n), \delta, \sigma = \Omega(1)\right]$

proof sketch: Last time we showed via query bundling and robustification that:

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

The starting PCP for QESAT$(\mathbb{F}_2)$ is $(\varepsilon_{PCP}, \delta_{PCP}, \text{Enc})$-sound (here Enc is the low-degree extension), and so is the resulting robust PCP for QESAT$(\mathbb{F}_2)$.

Hence it suffices to augment this latter with an encoding consistency test that is robust.

We know (from the robust PCP accepting wp $> \varepsilon_{PCP}$) that $\tilde{\pi}_a$ is δ_{PCP}-close to Enc$(\tilde{w})$ for some $\tilde{w} \in R[x]$.

We apply a "bespoke" query bundling and robustification to the prior slide's local decoder: the prover provides, $\forall i \in [|w|]$ and $r \in \mathbb{F}$, an encoding $e_{i,r}$ under a good code C of the coefficients of the polynomial $\hat{a}_{i,r}(z) := \text{Enc}(\tilde{w})(rz + i(1-z))$. The (relaxed) local decoder works as follows:

$$D^{\tilde{\pi}}(i) := \text{Sample } r_1, \dots, r_t \in \mathbb{F} \text{ and return plurality}_{j \in [t]} \{\hat{a}_{i,r_j}(0) \text{ where } \hat{a}_{i,r_j} := C^{-1}(e_{i,r_j})\},$$

provided that for random $\gamma \in \mathbb{F}$ $\hat{a}_{i,r_j}(\gamma) = \tilde{\pi}_a(r_j\gamma + i(1-\gamma))$.

PCP Theorem via Proof Composition

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

[everywhere below $\varepsilon_c = 0, \varepsilon_s = 1/2, \Sigma = \{0,1\}$]

Proof attempt Apply (non-interactive) proof composition theorem with:

- outer PCP: robust variant of the poly-size polylog-query PCP for NP [from last lecture]

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

- inner PCPP: proximity variant of the exp-size constant-query PCP for NP [from today]

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

By ensuring that $\sigma_{out} \geq \delta_{in}$ and setting $n_{in} = s_{out}(n)$, we get a composed PCP for NP with:

$CSAT \in PCP[\ell = \ell_{out} + 2^{r_{out}} \ell_{in} = exp(poly(\log n)) = n^{polylog(n)}, q = q_{in} = O(1), r = r_{out} + r_{in} = polylog(n)]$.

This PCP is too long.

Idea First, compose poly-size polylog-query PCP with itself to get smaller state size.

Second, compose the result with exp-size constant-query PCP.

This requires us to use a robust PCP of proximity.

PCP Theorem via Proof Composition

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

Part I of proof Apply (non-interactive) proof composition theorem with:

- outer PCP : robust variant of the poly-size polylog-query PCP for NP [like in prior slide]

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

- inner PCPP : robust & proximity variant of the poly-size polylog-query PCP for NP [from today]

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

By ensuring that $\sigma_{out} \geq \delta_{in}$ and setting $n_{in} = s_{out}(n)$, we get a composed PCP for NP with:

$CSAT \in PCP[\ell = \ell_{out} + 2^{r_{out}} \ell_{in} = poly(n), q = q_{in} = poly(\log \log n), r = r_{out} + r_{in} = O(\log n), s = s_{in} = poly(\log \log n), \sigma_{in} = \Omega(1)]$

In the next composition the composed PCP will act as the outer PCP.

Hence: (i) we used the fact that if the inner PCPP is robust then so is the composed PCP

(ii) we must keep track of the state size for the composed PCP (it is $s = s_{in}(n_{in})$)

PCP Theorem via Proof Composition

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

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

- outer PCP : the poly-size polylog log-query robust PCP for NP obtained from 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 exp-size constant-query PCP for NP [from today]

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

By ensuring that $\sigma_{out} \geq \delta_{in}$ and setting $n_{in} = s_{out}(n)$, we get a composed PCP for NP with:

$CSAT \in PCP[\ell = \ell_{out} + 2^{r_{out}} \ell_{in} = poly(n), q = q_{in} = O(1), r = r_{out} + r_{in} = O(\log n)]$ ← our goal! ■

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

<u>proof</u>: Similar 2-step composition but, in the first composition, start from an outer PCP that is robust & a proximity proof. Both compositions preserve the fact that the outer PCP is a proximity proof (and the proximity parameter remains unaffected). ■