Lecture 21

Foundations of Probabilistic Proofs
Fall 2020
Alessandro Chiesa

Robust PCPs

Today we construct robust PCPs for NP, which are used as an "outer PCP" in the proof of the PCP Theorem via proof composition.

Recall the definition of robust PCPs:

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

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

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

The above is for non-adaptive PCPs $\left(V^\pi(x;\rho)=D(x,\rho,\pi[Q(x,\rho)])\right)$ and $R(V)=\{((x,\rho),a) \mid a \in \Sigma^{Q(x,\rho)} \wedge D(x,\rho,a)=1\}$.

Robustness $\sigma \in [0,1/q)$ (wrt Hamming distance over Σ) is trivial.

The challenge is to achieve $\sigma = \Omega(1)$ even if q is super-constant.

We show how to achieve a robust analogue of the polynomial-size PCPs we constructed:

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

Proof Plan

We prove the theorem in two steps, starting from "standard" PCPs for NP.

[The construction for quadratic equations that uses the sumcheck protocol and low-degree test.]

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

|

Step 1: query bundling reduce query complexity to constant at the expense of alphabet size

↓

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

|

Step 2: robustification achieve constant robustness (over $\{0,1\}$) at the expense of query complexity

↓

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

We study each step in turn.

Robustification [1/3]

We wish to achieve good robustness over the binary alphabet, starting from a large-alphabet PCP.

Idea: break each large-symbol query into multiple bit queries

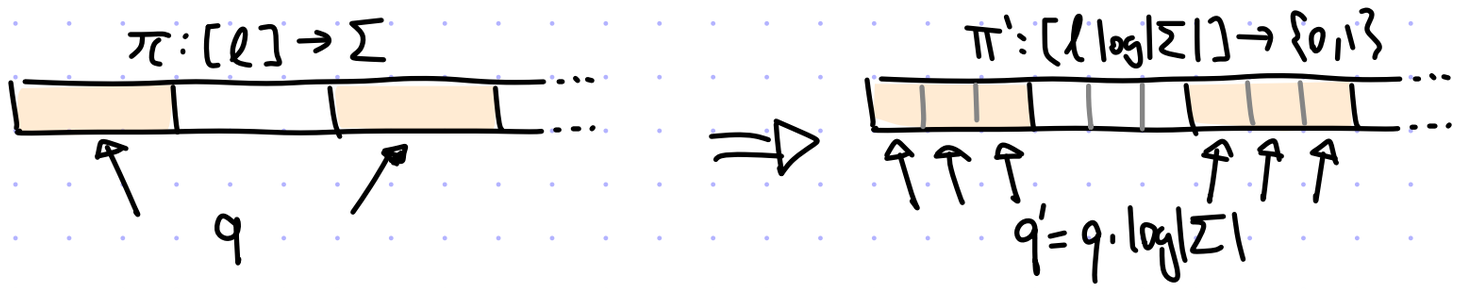

This preserves completeness and soundness, and indeed reduces the alphabet to binary.

Problem: the resulting PCP may have trivial robustness $\sigma \in [0, 1/q \cdot \log|\Sigma|)$

This is because many local views in the large-alphabet PCP could be 1 symbol (out of q) away from accepting, and so in the binary-alphabet PCP this could translate to 1 bit (out of $q \cdot \log|\Sigma|$) away from accepting.

This simple idea can be fixed to achieve this lemma:

no dependence on Σ

lemma: $PCP[\varepsilon_s, \Sigma, \ell, q, r] \subseteq PCP[\varepsilon_s, \Sigma' = \{0,1\}, \ell' = O(\ell \cdot \log|\Sigma|), q' = O(q \log|\Sigma|), r' = r, \sigma' = \Omega(\frac{1}{q})]$

Robustification [2/3]

lemma: $PCP[\varepsilon_s, \Sigma, \ell, q, r] \subseteq PCP[\varepsilon_s, \Sigma'=\{0,1\}, \ell'=O(\ell\cdot\log|\Sigma|), q'=O(q\log|\Sigma|), t'=r, \sigma'=\Omega(\frac{1}{q})]$

$\ell' = \ell \cdot N$ $\quad q' = q \cdot N$ $\quad \sigma' = \frac{\delta}{2q}$

proof: Encode each proof symbol via a good code to "sparsify" accepting local views.

Let $Enc: \Sigma \to \{0,1\}^N$ be an injective map with linear block length N (i.e., $O(\log|\Sigma|)$) and constant relative distance δ ($\Delta(Enc(u), Enc(u')) \geq \delta \geq \Omega(1)$ for every two distinct $u, u' \in \Sigma$).

This is a standard tool knows as a good code (it can be constructed via code concatenation).

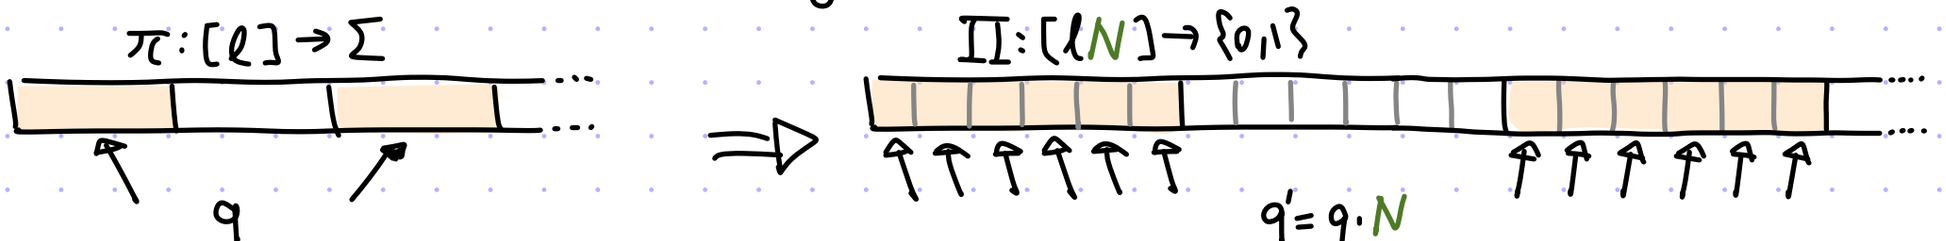

$P_r(x)$

1. $\pi := P(x) \in \Sigma^\ell$
2. For $i \in [\ell]$: $\Pi_i = Enc(\pi_i) \in \{0,1\}^N$
3. Output $\Pi = (\Pi_1, \ldots, \Pi_\ell) \in \{0,1\}^{N\cdot\ell}$

$V_r^{\Pi}(x)$

Run $V(x)$ by answering a query $i \in [\ell]$ by returning $Enc^{-1}(\Pi_i)$. (Reject if any $\perp$.)

Completeness: If $x \in L$ then $V(x)$ accepts $\pi = P(x)$ w.p. at least $1-\varepsilon_c$. Hence $V_r(x)$ accepts $\Pi := (Enc(\pi_1), \ldots, Enc(\pi_\ell))$ w.p. at least $1-\varepsilon_c$ because query $i \in [\ell]$ is answered with $Enc^{-1}(Enc(\pi_i)) = \pi_i$.

Robust soundness: next slide.

Robustification [3/3]

Suppose that $x \notin L$ and fix $\widetilde{\Pi} = (\widetilde{\Pi}_1, \ldots, \widetilde{\Pi}_\ell) \in \{0,1\}^{N \cdot \ell}$.

Let E be the event that the local view contains a string that is $\delta/2$-far from the code.

Then $\Pr_\rho\left[\Delta\left(\widetilde{\Pi}[Q_r(x,\rho)], R(V_r)[(x,\rho)]\right) \le \frac{\delta}{2 \cdot q}\right]$ (i.e. robustness parameter $\sigma = \delta/2q$)

$$\le \Pr_\rho\left[\Delta\left(\widetilde{\Pi}[Q_r(x,\rho)], R(V_r)[(x,\rho)]\right) \le \frac{\delta}{2 \cdot q} \,\middle|\, E\right] + \Pr\left[\Delta\left(\widetilde{\Pi}[Q_r(x,\rho)], R(V_r)[(x,\rho)]\right) \le \frac{\delta}{2 \cdot q} \,\middle|\, \bar{E}\right]$$

$$\le \underset{(a)}{0} + \underset{(b)}{\varepsilon_s}.$$

(a) If $\widetilde{\Pi}[Q_r(x,\rho)]$ contains a string that is $\delta/2$-far from the code then $\widetilde{\Pi}[Q_r(x,\rho)]$ is $\frac{\delta}{2 \cdot q}$-far from any accepting local view because all accepting local views consist of q strings in the code.

(b) The "correction" of $\widetilde{\Pi}$ is the string $\overline{\Pi} = (\overline{\Pi}_1, \ldots, \overline{\Pi}_\ell)$ where $\overline{\Pi}_i$ is closest codeword to $\widetilde{\Pi}_i \in \{0,1\}^N$ (breaking ties arbitrarily) and its decoding is $\overline{\pi} = (\overline{\pi}_1, \ldots, \overline{\pi}_\ell) := (\mathrm{Enc}^{-1}(\overline{\Pi}_1), \ldots, \mathrm{Enc}^{-1}(\overline{\Pi}_\ell)) \in \Sigma^\ell$.

By the soundness of V: $\Pr_\rho\left[V_r^{\overline{\Pi}}(x) = 1\right] = \Pr_\rho\left[V^{\overline{\pi}}(x) = 1\right] \le \varepsilon_s$

Whenever every string in $\widetilde{\Pi}[Q_r(x,\rho)]$ is $\delta/2$-close to the code then $\overline{\Pi}[Q_r(x,\rho)]$ is the only string of q codewords that is $\frac{\delta}{2 \cdot q}$-close to $\widetilde{\Pi}[Q_r(x,\rho)]$. So

$$\Pr\left[\Delta\left(\widetilde{\Pi}[Q_r(x,\rho)], R(V_r)[(x,\rho)]\right) \le \frac{\delta}{2 \cdot q} \,\middle|\, \bar{E}\right] \le \Pr_\rho\left[\overline{\Pi}[Q_r(x,\rho)] \in R(V_r)[(x,\rho)]\right] = \Pr_\rho\left[V_r^{\overline{\Pi}}(x) = 1\right]. \quad \blacksquare$$

Bundling Queries [1/4]

We are given a PCP system over a small alphabet, e.g., $\Sigma=\{0,1\}$.

Goal: a new PCP system with $O(1)$ queries over a larger alphabet, e.g., $\Sigma'=\Sigma^{\approx q}$.

We restrict our attention to non-adaptive PCPs: $V^{\pi}(x;\rho)=D(x,\rho,\pi[Q(x,\rho)])$

Idea: write answers to query sets as large symbols and add a consistency test

The new PCP is $\pi_b := \Big($ 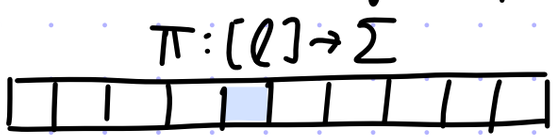

$,\ \Big($ 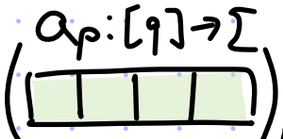

$\Big)_{\rho\in\{0,1\}^r}\Big)$

The new PCP verifier $V_b(x)$ is:

1. Sample $\rho\in\{0,1\}^r$ and $i\in[q]$.
2. Read $a_\rho\in\Sigma^q$ and $\pi[Q(x,\rho)[i]]$.
3. Check that $a_\rho[i]=\pi[Q(x,\rho)[i]]$.
4. Check that $V(x;\rho)=1$ when answering j-th query with $a_\rho[j]$.

$a_\rho=\pi[Q(x,\rho)]$ in an honest proof

lemma: $PCP[\varepsilon_s,\Sigma,\ell,q,r]\subseteq PCP[\varepsilon_s'=1-(1-\varepsilon_s)\cdot\frac{1}{q},\Sigma'=\Sigma^q,\ell'=O(\ell+2^r),q'=2,r'=r+\log q]$

Does not suffice for us: we need constant soundness error even when q is super constant.

Bundling Queries [2/4]

Fix a field $\mathbb{F}$, subset $H \subseteq \mathbb{F}$, and number of variables $m \in \mathbb{N}$. Assume that $|\mathbb{F}| \geq \max\{|\Sigma|, q\}$. (for convenience)

We identify $[\ell]$ with H^m by setting $m := \log \ell / \log |H|$.

Similarly, we can view a query set $Q(x,\rho) \subseteq [\ell]$ as q elements of H^m.

Two changes from prior approach:

- replace $\pi : [\ell] \to \Sigma$ with its $(\mathbb{F}, H, m)$-extension $\hat{\pi} : \mathbb{F}^m \to \mathbb{F}$
- replace $a_\rho : [q] \to \Sigma$ with $\hat{a}_\rho(z) := \hat{\pi}(Q_\rho(z)) \in \mathbb{F}^{<q \cdot m \cdot |H|}[z]$ where $Q_\rho : \mathbb{F} \to \mathbb{F}^m$ is m polynomials of degree less than q s.t. $\forall\, j \in [q]$ $Q_\rho(j) = Q(x,\rho)[j] \in H^m$.

(Here we use $1, 2, \ldots, q$ to denote any q distinct elements in $\mathbb{F}$.)

Warm up: $\pi_b := \Big(\hat{\pi} : \mathbb{F}^m \to \mathbb{F}, \big(\hat{a}_\rho(z)\big)_{\rho \in \{0,1\}^r} \Big)$

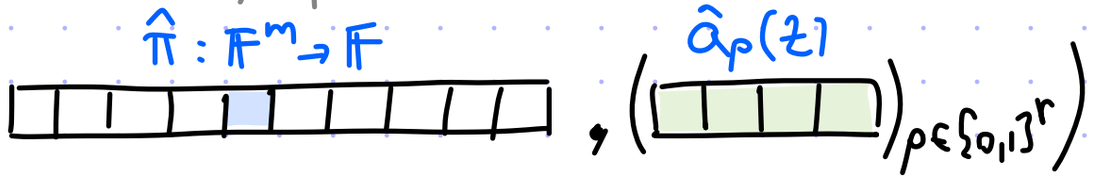

$P_b(x)$

1. $\pi := P(x) \in \mathbb{F}^{H^m}$
2. $\hat{\pi} : \mathbb{F}^m \to \mathbb{F}$ is $(\mathbb{F}, H, m)$-extension of π
3. For $\rho \in \{0,1\}^r$: $\hat{a}_\rho(z) := \hat{\pi}(Q_\rho(z))$
4. Output $\pi_b := (\hat{\pi}, (\hat{a}_\rho)_{\rho \in \{0,1\}^r})$

$V_b(x)$

1. Sample $\rho \in \{0,1\}^r$ and $\gamma \in \mathbb{F}$.
2. Read $\hat{a}_\rho \in \mathbb{F}[z]$ and $\hat{\pi}(Q_\rho(\gamma))$.
3. Check that $\hat{a}_\rho(\gamma) \stackrel{?}{=} \hat{\pi}(Q_\rho(\gamma))$.
4. Check that $V(x;\rho) = 1$ when answering j-th query with $\hat{a}_\rho(j)$.

Bundling Queries [3/4]

$\pi_b := \Big($  $,$ 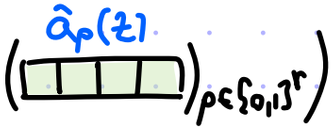 $\Big)_{\rho \in \{0,1\}^r}$

$V_b(x)$

1. Sample $\rho \in \{0,1\}^r$ and $\gamma \in \mathbb{F}$.
2. Read $\hat{a}_\rho \in \mathbb{F}[z]$ and $\hat{\pi}(Q_\rho(\gamma))$.
3. Check that $\hat{a}_\rho(\gamma) \stackrel{?}{=} \hat{\pi}(Q_\rho(\gamma))$.
4. Check that $V(x;\rho)=1$ when answering j-th query with $\hat{a}_\rho(j)$.

For this warm-up case ($\hat{\pi}$ is guaranteed to be the low degree extension), we can prove:

claim: the soundness error is at most $1-(1-\varepsilon_s)\cdot\left(1-\frac{qm|H|}{|\mathbb{F}|}\right)$. [improving on $1-(1-\varepsilon_s)\cdot\frac{1}{q}$]

proof: Suppose that $x \notin L$ and fix a PCP $\pi_b = (\hat{\pi}, (\hat{a}_\rho)_{\rho\in\{0,1\}^r})$.

$$\Pr\left[V_b^{\pi_b}(x)=1\right] = \Pr_{\rho,\gamma}\left[D(x,\rho,\hat{a}_\rho(1),\dots,\hat{a}_\rho(q))=1 \wedge \hat{a}_\rho(\gamma)=\hat{\pi}(Q_\rho(\gamma))\right]$$

$$= 1-\Pr_{\rho,\gamma}\left[D(x,\rho,\hat{a}_\rho(1),\dots,\hat{a}_\rho(q))=0 \vee \hat{a}_\rho(\gamma)\neq\hat{\pi}(Q_\rho(\gamma))\right] \leq 1-(1-\varepsilon_s)\cdot\left(1-\frac{qm|H|}{|\mathbb{F}|}\right) \quad \blacksquare$$

Problem: how to handle the case where $\hat{\pi}$ is noisy?

The LDT that we say makes $\omega(1)$ queries to f. Moreover, the query $Q_\rho(\gamma)$ to f is biased.

Bundling Queries [4/4]

- We rely on a large-alphabet constant-query low-degree test:

  theorem [line vs. point test] $\forall f: \mathbb{F}^m \to \mathbb{F}$ $\forall$ lines oracle $\{\hat{g}_{a,b}\}_{a,b\in\mathbb{F}^m}$ of degree d
  if f is δ-far from total degree d then $\Pr_{a,b\in\mathbb{F}^m,\mu\in\mathbb{F}}[f(a\mu+b)=\hat{g}_{a,b}(\mu)] \le \varepsilon_{LST}(\delta)$

- We randomize the query to $\hat{\pi}$: $Q_{\rho,v}(z)$ is the polynomial s.t. $Q_{\rho,v}(q+1)=v$ for random $v\in\mathbb{F}^m$, so that $\forall \gamma\in\mathbb{F}\setminus[q]$ $Q_{\rho,v}(\gamma)$ is random in $\mathbb{F}^m$

Here is the final construction:

$P_b(x)$
1. $\pi := P(x) \in \mathbb{F}^{H^m}$
2. $\hat{\pi}: \mathbb{F}^m \to \mathbb{F}$ is $(\mathbb{F},H,m)$-extension of π
3. For $\rho\in\{0,1\}^r$ & $v\in\mathbb{F}^m$: $\hat{a}_{\rho v}(z) := \hat{\pi}(Q_{\rho v}(z))$
4. For $a,b\in\mathbb{F}^m$: $\hat{g}(z) := \hat{\pi}(az+b)$
5. Output $\pi_b := (\hat{\pi}, (\hat{a}_{\rho v})_{\rho\in\{0,1\}^r, v\in\mathbb{F}^m}, (\hat{g}_{a,b})_{a,b\in\mathbb{F}^m})$

$\pi_b := \left(\hat{\pi}:\mathbb{F}^m\to\mathbb{F}, \left(\hat{a}_{\rho v}(z)\right)_{\rho\in\{0,1\}^r, v\in\mathbb{F}^m}, \left(\hat{g}_{a,b}(z)\right)_{a,b\in\mathbb{F}^m}\right)$ (lines oracle of degree $d=m\cdot|H|$)

$V_b(x)$
1. Sample $a,b\in\mathbb{F}^m$ & $\mu\in\mathbb{F}$, and check $\hat{g}_{a,b}(\mu) \stackrel{?}{=} \hat{\pi}(a\mu+b)$.
2. Sample $\rho\in\{0,1\}^r$, $v\in\mathbb{F}^m$, $\gamma\in\mathbb{F}\setminus[q]$, and check $\hat{a}_{\rho,v}(\gamma) \stackrel{?}{=} \hat{\pi}(Q_{\rho v}(\gamma))$.
3. Check that $V(x;\rho)=1$ when answering j-th query with $\hat{a}_\rho(j)$.

One can prove that the new error is $\varepsilon_\delta' = \max\left\{\varepsilon_{LST}(\delta), 1-(1-\varepsilon_\delta)\left(1-\frac{qm|H|}{|\mathbb{F}|-q}-\delta\right)\right\}$.

lemma: $PCP[\varepsilon_\delta,\Sigma,\ell,q,r] \subseteq PCP[\varepsilon_\delta', \Sigma'=\Sigma^{qm|H|}, \ell'=|\mathbb{F}|^m+2^r|\mathbb{F}|+|\mathbb{F}|^{2m}, q'=O(1), r'=r+O(\log|\mathbb{F}|)]$

We can then set $|H|=\log\ell$, $m=\frac{\log\ell}{\log|H|}=\frac{\log\ell}{\log\log\ell}$, $|\mathbb{F}|=O(qm|H|)$ for the regime of interest.