Foundations of Probabilistic Proofs

A course by **Alessandro Chiesa**

Lecture 20

Robust PCPs

Robust PCPs

We construct robust PCPs for NP.

They are used as "outer PCP" in the proof of the PCP Theorem via proof composition.

We consider non-adaptive verifiers: $V^{\pi}(x;s) = D(S(x,\rho), \pi[Q(x,s)])$.

decision algorithm · state algorithm · query algorithm

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

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: $\forall x \notin L(R) \quad \forall \tilde{\pi} \quad \Pr_s\left[\Delta(\tilde{\pi}[Q(x,s)], R(V)[S(x,s)]) \leq \sigma\right] \leq \varepsilon_s$.

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 achieve a robust analogue of the poly-length polylog-query PCP:

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

Proof Plan

We prove the theorem in two steps, starting from the "canonical" PCP for NP.

$NP \subseteq PCP\left[\varepsilon_c = 0,\ \varepsilon_s = \frac{1}{2},\ \Sigma = \{0,1\},\ \ell = \text{poly}(n),\ q = \text{poly}(\log n),\ r = O(\log n)\right]$ ← PCP for QESAT that uses the sumcheck protocol and the low-degree test

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

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

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

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

We study each step.

Robustification [1/4]

GOAL: achieve good robustness over the binary alphabet, starting from a large-alphabet PCP.

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

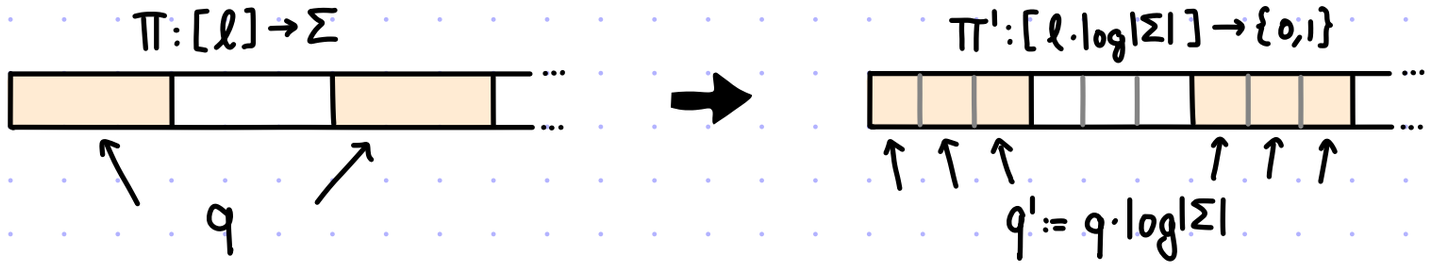

This preserves completeness and soundness, and reduces the alphabet to binary.

PROBLEM: the resulting PCP may have trivial robustness $\sigma \in [0, \frac{1}{q \cdot \log|\Sigma|})$.

Many local views in the large-alphabet PCP may be 1 symbol (out of q) away from accepting.

In the binary-alphabet PCP, each such view may be 1 bit (out of $q \cdot \log|\Sigma|$) away from accepting.

The simple idea can be fixed to achieve this lemma:

lemma: $PCP[\varepsilon_c, \varepsilon_s, \Sigma, \ell, q, r]$

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

no dependence on Σ

Robustification [2/4]

IDEA: "sparsify" accepting local views by encoding each proof symbol via an error-correcting code

Let $\mathsf{Enc}: \Sigma \to \{0,1\}^N$ be an injective map with relative distance δ ($\forall$ distinct $a,b \in \Sigma$ $\Delta(\mathsf{Enc}(a), \mathsf{Enc}(b)) \geq \delta$).

lemma: $\mathrm{PCP}[\varepsilon_c, \varepsilon_s, \Sigma, \ell, q, r] \subseteq \mathrm{PCP}[\varepsilon_c, \varepsilon_s, \Sigma'=\{0,1\}, \ell' = \ell \cdot N, q' = q \cdot N, r' = r, \delta = \frac{\delta}{4q}]$

The prior lemma follows from the fact that $\exists$ Enc with $N = O(\log|\Sigma|)$ and $\delta = \Omega(1)$.

standard code construction via code concatenation

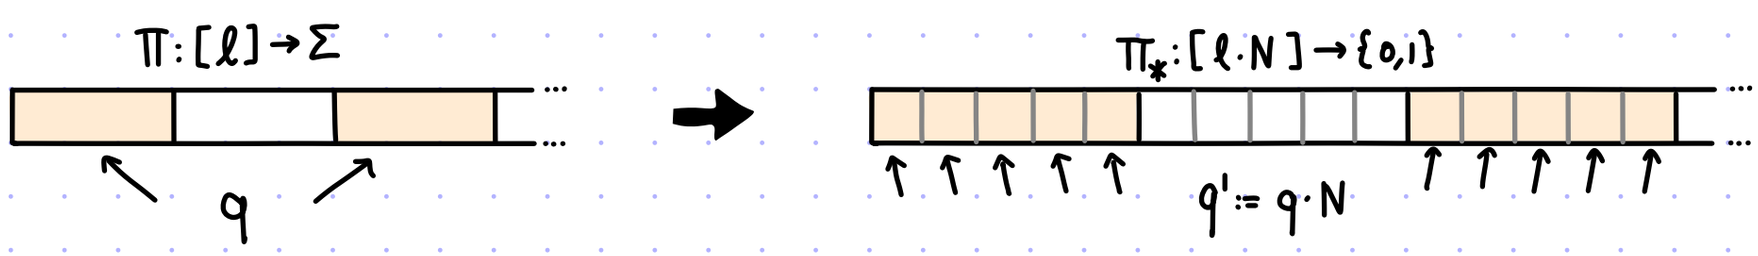

$P_*(x,w)$

1. $\pi := P(x,w) \in \Sigma^\ell$
2. $\forall i \in [\ell]: c_i := \mathsf{Enc}(\pi[i]) \in \{0,1\}^N$
3. Output $\pi_* := (c_i)_{i \in [\ell]} \in \{0,1\}^{N \cdot \ell}$

$V_*^{\pi_*}(x)$

Run $V(x)$ by answering each query $i \in [\ell]$:

- make N queries to read $c_i \in \{0,1\}^N$
- return $a_i := \mathsf{Enc}^{-1}(c_i) \in \Sigma$ (reject if $a_i = \perp$)

Completeness: If $(x,w) \in R$ then $\Pr[V^\pi(x) = 1 \mid \pi \leftarrow P(x,w)] \geq 1 - \varepsilon_c$. Since $P_*(x,w)$ outputs $\pi_* = (\mathsf{Enc}(\pi[i]))_{i \in [\ell]}$, $V_*(x)$ answers each query $i \in [\ell]$ of $V(x)$ with $\mathsf{Enc}^{-1}(\mathsf{Enc}(\pi[i])) = \pi[i]$.

Robustification [3/4]

Robust soundness: Fix $x \notin L(R)$ and $\widetilde{\pi}_* = (\widetilde{c}_i)_{i \in [\ell]} \in \{0,1\}^{N \cdot \ell}$.

Define $C := Enc(\Sigma) \subseteq \{0,1\}^N$ and the interleaved code $C^q := \{(c_1,\ldots,c_q) : \forall i \in [q]\ c_i \in C\} \subseteq \{0,1\}^{N \cdot q}$.

Observe that $\forall s \in \{0,1\}^r$ $R(V_*)[S_*(x,s)] \subseteq C^q$ (every accepting local view is a codeword in C^q).

The relative distance of C^q is $\geq \frac{\delta}{q}$ (since the relative distance of C is $\geq \delta$),

so the unique-decoding radius of C^q is $\frac{\delta}{2 \cdot q}$ ($\forall\, u \in \{0,1\}^{N \cdot q}$ there is at most one $c \in C^q$ s.t. $\Delta(u,c) < \frac{\delta}{2 \cdot q}$).

Define the event E = "local view $\widetilde{\pi}_*[Q_*(x,s)]$ contains $\widetilde{c}_i$ that is at (relative) distance $\geq \delta/2$ from C".

Then $\Pr_s\left[\Delta(\widetilde{\pi}_*[Q_*(x,s)], R(V_*)[S_*(x,s)]) < \frac{\delta}{2 \cdot q}\right]$ (any robustness parameter $\sigma < \frac{\delta}{2q}$, e.g. $\sigma = \frac{\delta}{4q}$)

$\leq \Pr_s\left[\Delta(\widetilde{\pi}_*[Q_*(x,s)], R(V_*)[S_*(x,s)]) < \frac{\delta}{2 \cdot q} \mid E\right] + \Pr_s\left[\Delta(\widetilde{\pi}_*[Q_*(x,s)], R(V_*)[S_*(x,s)]) < \frac{\delta}{2 \cdot q} \mid \overline{E}\right]$

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

⊛ The bound $\frac{\delta}{2q}$ can be improved to $\lceil qN \frac{\delta}{2q} \rceil / qN = \frac{\lceil N\delta/2 \rceil}{qN}$ ($\frac{\delta}{2q}$ rounded up to the next multiple of $\frac{1}{qN}$) via the same analysis.

(a) Suppose that E holds ($\widetilde{\pi}_*[Q_*(x,s)]$ contains $\widetilde{c}_i$ that is at relative distance $\geq \delta/2$ from C).

Then $\widetilde{\pi}_*[Q_*(x,s)]$ is at (relative) distance $\geq \delta/2q$ from C^q and thus from any accepting local view,

because every accepting local view consists of q strings in C (i.e. $R(V_*)[S_*(x,s)] \subseteq C^q$).

Robustification [4/4]

ⓑ We are left to prove that $\Pr_s\left[\Delta(\widetilde{\Pi}_*[Q_*(x,s)], R(V_*)[S_*(x,s)]) < \frac{\delta}{2\cdot q} \mid \overline{E}\right] \leq \varepsilon_s$.

Define the "correction" $\overline{\Pi}_* := (\bar{c}_i)_{i\in[\ell]}$ where $\bar{c}_i \in C$ is closest to $\tilde{c}_i$ (break ties arbitrarily).

Define its decoding $\pi := (\mathrm{Enc}^{-1}(\bar{c}_i))_{i\in[\ell]} \in \Sigma^\ell$.

By the soundness of V, $\Pr\left[V_*^{\overline{\Pi}_*}(x)\right] = \Pr\left[V^{\pi}(x)=1\right] \leq \varepsilon_s$.

If E does not hold (every string in $\widetilde{\Pi}_*[Q_*(x,s)]$ is at relative distance $< \delta/2$ to C), then no codeword in C^q is closer to $\widetilde{\Pi}_*[Q_*(x,s)]$ than $\overline{\Pi}_*[Q_*(x,s)]$ is.

Indeed, $\forall\, c \in C^q \quad \Delta(\widetilde{\Pi}_*[Q_*(x,s)], c) = \frac{1}{q}\sum_{j=1}^{q} \Delta(\tilde{c}_{Q(x,s)[j]}, c_j)$

$\geq \frac{1}{q}\sum_{j=1}^{q} \Delta(\tilde{c}_{Q(x,s)[j]}, \bar{c}_{Q(x,s)[j]}) = \Delta(\widetilde{\Pi}_*[Q_*(x,s)], \overline{\Pi}_*[Q_*(x,s)])$.

$\forall i\in[\ell]$ $\bar{c}_i \in C$ is closest to $\tilde{c}_i$

Hence $\overline{\Pi}_*[Q_*(x,s)]$ is the ONLY local view that can satisfy $\Delta(\widetilde{\Pi}_*[Q_*(x,s)], \overline{\Pi}_*[Q_*(x,s)]) < \frac{\delta}{2\cdot q}$. (unique-decoding radius of C^q)

Since $R(V_*)[S_*(x,s)] \subseteq C^q$, if $\Delta(\widetilde{\Pi}_*[Q_*(x,s)], R(V_*)[S_*(x,s)]) < \frac{\delta}{2\cdot q}$

then $\Delta(\widetilde{\Pi}_*[Q_*(x,s)], \overline{\Pi}_*[Q_*(x,s)]) < \frac{\delta}{2\cdot q}$ and $\overline{\Pi}_*[Q_*(x,s)] \in R(V_*)[S_*(x,s)]$.

We conclude that

$$\Pr_s\left[\Delta(\widetilde{\Pi}_*[Q_*(x,s)], R(V_*)[S_*(x,s)]) < \frac{\delta}{2\cdot q} \mid \overline{E}\right] \leq \Pr_s\left[\overline{\Pi}_*[Q_*(x,s)] \in R(V_*)[S_*(x,s)]\right] = \Pr\left[V_*^{\overline{\Pi}_*}(x)=1\right]. \quad \blacksquare$$

Bundling Queries

GOAL: reduce query complexity to constant at the expense of alphabet size

IDEA: provide the answer to each query set (as a large symbol) + consistency test

$P_*(x,w)$

1. $\pi := P(x,w) \in \Sigma^\ell$.
2. For every $s \in \{0,1\}^r$:
   $a_s := \pi[Q(x,s)] \in \Sigma^q$.
3. Output $\pi_* := (\pi, (a_s)_{s\in\{0,1\}^r})$.

$V_*^{\pi_*}(x)$

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

$\pi_* := \Big($ 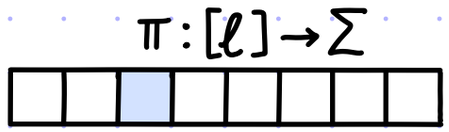 $,$ 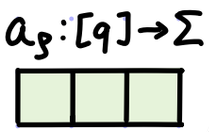 $\Big)_{s\in\{0,1\}^r}$

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

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

Nevertheless, we prove soundness because the analysis is a useful warm-up.

Bundling Queries

lemma: $\mathrm{PCP}[\varepsilon_c, \varepsilon_s, \Sigma, \ell, q, r] \leq \mathrm{PCP}[\varepsilon_c, \varepsilon_s' = 1 - \frac{1-\varepsilon_s}{q}, \Sigma' = \Sigma^q, \ell' = \ell + 2^r, q' = 2, r' = r + \log q]$

Proof of soundness.

Fix $x \notin L(R)$ and $\pi_* = (\pi, (a_s)_{s \in \{0,1\}^r})$.

$$\Pr[V_*^{\pi_*}(x) = 1] = \Pr[V_*^{\pi_*}(x) = 1 \mid V^{\pi}(x) = 1] \cdot \Pr[V^{\pi}(x) = 1] + \Pr[V_*^{\pi_*}(x) = 1 \mid V^{\pi}(x) = 0] \cdot \Pr[V^{\pi}(x) = 0]$$

$$\leq 1 \cdot \Pr[V^{\pi}(x) = 1] + \Pr[V_*^{\pi_*}(x) = 1 \mid V^{\pi}(x) = 0] \cdot (1 - \Pr[V^{\pi}(x) = 1])$$

$$\leq 1 \cdot \Pr[V^{\pi}(x) = 1] + (1 - \tfrac{1}{q}) \cdot (1 - \Pr[V^{\pi}(x) = 1])$$

$$= \tfrac{1}{q} \cdot \Pr[V^{\pi}(x) = 1] + 1 - \tfrac{1}{q} \leq \tfrac{1}{q} \cdot \varepsilon_s + 1 - \tfrac{1}{q} = 1 - \tfrac{1-\varepsilon_s}{q}.$$

We are left to show the inequality highlighted in green:

$$\Pr[V_*^{\pi_*}(x) = 1 \mid V^{\pi}(x) = 0] \leq \Pr_{s,\gamma}\left[V_*^{\pi_*}(x; (s,\gamma)) = 1 \;\middle|\; \begin{matrix} a_s \neq \pi[Q(x,s)] \\ \vee\; V^{[Q(x,s), a_s]}(x; s) = 0 \end{matrix} \right]$$

$$\leq \Pr_{s,\gamma}\left[V_*^{\pi_*}(x; (s,\gamma)) = 1 \;\middle|\; a_s \neq \pi[Q(x,s)] \right]$$

$$\leq \Pr_{s,\gamma}\left[a_s[\gamma] = \pi[Q(x,s)][\gamma] \;\middle|\; a_s \neq \pi[Q(x,s)] \right]$$

$$\leq 1 - \tfrac{1}{q}.$$

■

Bundling Queries [3/7]

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)

Identify $[\ell]$ with H^m by setting $m := \frac{\log \ell}{\log |H|}$. 

View a query set $Q(x,s) \subseteq [\ell]$ as q elements in H^m.

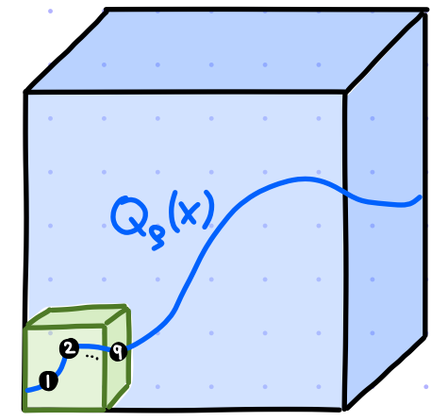

Changes from prior approach:

- replace $\pi : [\ell] \to \Sigma$ with its $(\mathbb{F}, H, m)$-extension $\hat{\pi} : \mathbb{F}^m \to \mathbb{F}$
- replace $a_s : [q] \to \Sigma$ with $\hat{a}_s(z) := \hat{\pi}(Q_s(z)) \in \mathbb{F}^{<q \cdot m \cdot |H|}[z]$ where

  $Q_s : \mathbb{F} \to \mathbb{F}^m$ is m polynomials of degree $<q$ s.t. $\forall j \in [q]$ $Q_s(j) := Q(x,s)[j] \in H^m$.

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

Warm Up: $\pi_* := \left(\hat{\pi} : \mathbb{F}^m \to \mathbb{F}, \left(\hat{a}_s(z) \right)_{s \in \{0,1\}^r} \right)$

$P_*(x,w)$

1. $\pi := P(x,w) \in \Sigma^\ell \subseteq \mathbb{F}^{H^m}$.
2. $\hat{\pi} : \mathbb{F}^m \to \mathbb{F}$ is $(\mathbb{F}, H, m)$-extension of π.
3. $\forall s \in \{0,1\}^r$: $\hat{a}_s(z) := \hat{\pi}(Q_s(z))$.
4. Output $\pi_* := (\hat{\pi}, (\hat{a}_s)_{s \in \{0,1\}^r})$.

$V_*^{\pi_*}(x)$

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

Bundling Queries

$\pi_* = \Big($ 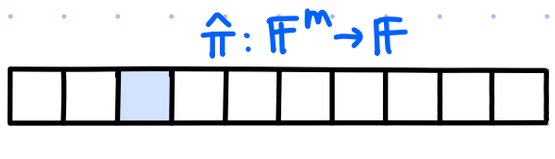

$,\ \Big($ 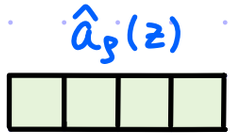

$\Big)_{s \in \{0,1\}^r} \Big)$

$V_*^{\pi_*}(x)$

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

For this warm-up case, we assume that $\hat{\pi}$ is an $(\mathbb{F}, H, m)$-extension of (some) $\pi : [\ell] \to \Sigma$.

claim: the soundness error is $\leq 1 - (1-\varepsilon_s) \cdot \left(1 - \frac{q \cdot m \cdot |H|}{|\mathbb{F}|}\right)$. [this improves on $1 - (1-\varepsilon_s) \cdot \frac{1}{q}$]

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

It suffices to show that $\Pr[V_*^{\pi_*}(x)=1 \mid V^{\pi}(x)=0] \leq \frac{q \cdot m \cdot |H|}{|\mathbb{F}|}$ (by a similar analysis as the prior construction).

$$\begin{aligned}
\Pr[V_*^{\pi_*}(x)=1 \mid V^{\pi}(x)=0] &\leq \Pr_{s,\gamma}\left[V_*^{\pi_*}(x;(s,\gamma))=1 \;\middle|\; \begin{matrix} \hat{a}_s \neq \hat{\pi}[Q_s] \\ \vee\ V^{[Q_s([q]),\hat{a}_s([q])]}(x;s)=0 \end{matrix} \right] \\
&\leq \Pr_{s,\gamma}\left[V_*^{\pi_*}(x;(s,\gamma))=1 \;\middle|\; \hat{a}_s \neq \hat{\pi}[Q_s] \right] \\
&\leq \Pr_{s,\gamma}\left[\hat{a}_s[\gamma] = \hat{\pi}[Q_s(\gamma)] \;\middle|\; \hat{a}_s \neq \hat{\pi}[Q_s] \right] \\
&\leq \frac{q \cdot m \cdot |H|}{|\mathbb{F}|}.
\end{aligned}$$

■

Bundling Queries

Q: how to handle the noisy case?

That is $\pi_* = (f, (\hat{a}_s)_{s\in\{0,1\}^r})$ where $f: \mathbb{F}^m \to \mathbb{F}$ is arbitrary.

We no longer require that f is the $(\mathbb{F},H,m)$-extension $\hat{\pi}$ of some $\pi:[\ell]\to\Sigma$.

PROBLEM 1: The Rubinfeld-Sudan LDT that we studied makes $\omega(1)$ queries to f.

In fact, every LDT makes $d+2 = \omega(1)$ queries to f.

FIX 1: We use a large-alphabet constant-query PCPP for low-degreeness.

The LINE VS. POINT TEST is such a PCPP.

$$P_{LPT}(f) := \left(\hat{g}_{a,b} := f(ax+b) \in \mathbb{F}^{\le d}[x]\right)_{a,b\in\mathbb{F}^m}$$

$V_{LPT}^{f,(\hat{g}_{a,b})_{a,b\in\mathbb{F}^m}}$

1. Sample $a,b \in \mathbb{F}^m$ and $\mu \in \mathbb{F}$.
2. Check that $f(a\mu+b) = \hat{g}_{a,b}(\mu)$.

Completeness: $\forall f:\mathbb{F}^m\to\mathbb{F}$ of total degree $\le d$ for $\pi_{px} := P_{LPT}(f)$, $\Pr\left[V_{LPT}^{f,\pi_{px}} = 1\right] = 1$.

Soundness: $\forall f:\mathbb{F}^m\to\mathbb{F}$ $\forall \tilde{\pi}_{px} = (\hat{g}_{a,b} \in \mathbb{F}^{\le d}[x])$, f is δ-far from total degree $d \to \Pr\left[V_{LPT}^{f,\tilde{\pi}_{px}} = 1\right] \le \varepsilon_{LPT}(\delta)$.

Fact [that we do not prove]: $\exists\, \delta_0, \varepsilon_0 \in (0,1)$ s.t. $\delta \ge \delta_0 \to \varepsilon_{LPT}(\delta) \le \varepsilon_0$.

Bundling Queries [6/7]

Q: how to handle the noisy case?

That is $\pi_* = (f, (\hat{a}_s)_{s\in\{0,1\}^r})$ where $f: \mathbb{F}^m \to \mathbb{F}$ is arbitrary. ← We no longer require that f is the $(\mathbb{F},H,m)$-extension $\hat{\pi}$ of some $\pi:[\ell]\to\Sigma$.

PROBLEM 2: The query $Q_s(\gamma)$ to f is NOT uniformly random in $\mathbb{F}^m$.

Indeed, $$Q_s(\gamma) := \sum_{j\in[q]} Q(x,s)[j]\cdot L_{[q],j}(\gamma)$$

where $\{L_{[q],j}(x)\}_{j\in[q]}$ are the Lagrange polynomials for $[q]$.

The degree of each $L_{[q],j}(x)$ is $q-1$, and $L_{[q],j}(\gamma)$ is (typically) NOT uniformly random for random $\gamma\leftarrow\mathbb{F}$.

E.g. the distribution $\{\gamma^2 \mid \gamma\leftarrow\mathbb{F}\}$ is supported only on squares of $\mathbb{F}$. (Approx half the elements if $\mathrm{char}(\mathbb{F})\neq 2$.)

Fix 2: We randomize the query to f.

For every $v\in\mathbb{F}^m$,

$Q_{sv}: \mathbb{F}\to\mathbb{F}^m$ is m polynomials of degree $\le q$ s.t. $\begin{cases} \forall j\in[q]\ Q_{sv}(j) := Q(x,s)[j] \in H^m \\ Q_{sv}(q+1) := v \end{cases}$.

Note that, for random $v\in\mathbb{F}^m$, $\forall\ \gamma\in\mathbb{F}\setminus[q]$ $Q_{sv}(\gamma)$ is random in $\mathbb{F}^m$.

Indeed, $$Q_s(\gamma) := \sum_{j\in[q]} Q(x,s)[j]\cdot L_{[q+1],j}(\gamma) + v\cdot L_{[q+1],q+1}(\gamma)$$

and $L_{[q+1],q+1}(\gamma)\neq 0$ for $\gamma\in\mathbb{F}\setminus[q]$.

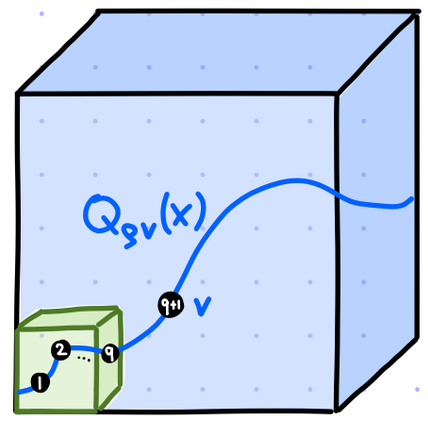

Bundling Queries

The final construction.

lines oracle of degree $d < m \cdot |H|$

$$\pi_* = \Big(\underbrace{\square\square\square\square\square\square\square\square\square}_{f: \mathbb{F}^m \to \mathbb{F}},\ \big(\underbrace{\square\square\square\square}_{\hat{a}_{s,v}(z)} \big)_{\substack{s \in \{0,1\}^r \\ v \in \mathbb{F}^m}},\ \big(\underbrace{\square\square}_{\hat{g}_{a,b}} \big)_{a,b \in \mathbb{F}^m} \Big)$$

$P_*(x,w)$

1. $\pi := P(x,w) \in \Sigma^\ell \subseteq \mathbb{F}^{H^m}$.
2. $\hat{\pi}: \mathbb{F}^m \to \mathbb{F}$ is $(\mathbb{F},H,m)$-extension of π.
3. $\forall\, s \in \{0,1\}^r,\ v \in \mathbb{F}^m$: $\hat{a}_{s,v}(z) := \hat{\pi}(Q_{s,v}(z))$.
4. $\forall\, a,b \in \mathbb{F}^m$: $\hat{g}_{a,b}(z) := \hat{\pi}(az+b)$.
5. Output $\pi_* := (\hat{\pi}, (\hat{a}_{s,v})_{\substack{s \in \{0,1\}^r \\ v \in \mathbb{F}^m}}, (\hat{g}_{a,b})_{a,b \in \mathbb{F}^m})$.

$V_*^{\pi_*}(x)$

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

lemma: $$PCP[\varepsilon_c, \varepsilon_s, \Sigma, \ell, q, r] \subseteq PCP\left[\begin{array}{ll} \varepsilon_c & \varepsilon_s' = \max\left\{\varepsilon_{LPT}(\delta),\ 1-(1-\varepsilon_s)\cdot\left(1-\frac{qm|H|}{|\mathbb{F}|-q}-\delta\right)\right\} \\ \Sigma' = \Sigma^{qm|H|} & \ell' = |\mathbb{F}|^m + 2^r|\mathbb{F}|^m + |\mathbb{F}|^{2m} \\ q' = 4 & r' = r + O(m \cdot \log|\mathbb{F}|) \end{array}\right]$$

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

Bibliography

Robust PCPs

- [DR 2006]: Assignment testers: towards a combinatorial proof of the PCP theorem, by Irit Dinur, Omer Reingold.
- [ALMSS 1998]: Proof verification and the hardness of approximation problems, by Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, Mario Szegedy.
- [BGHSV 2005]: Robust PCPs of proximity, shorter PCPs and applications to coding, by Eli Ben-Sasson, Oded Goldreich, Prahladh Harsha, Madhu Sudan, Salil Vadan.

Section 7.2 on O(1) query tests