Lecture B.5

Analysis of FRI

(Fast Reed-Solomon IOP)

Tom Gur

Summer Graduate School on
Foundations and Frontiers of Probabilistic Proofs
July 30 2021

Recap

Proximity Proofs for the Reed--Solomon Code

We say that (P,V) is an IOP of proximity (IOPP) for $RS[\mathbb{F},L,d]$ if:

① <u>completeness</u>: if $f \in RS[\mathbb{F},L,d]$ then $\Pr[\langle P(f), V^f\rangle = 1] = 1$
② <u>soundness</u>: if f is δ-far from $RS[\mathbb{F},L,d]$ then $\forall \tilde{P}$ $\Pr[\langle \tilde{P}, V^f\rangle = 1] \le \varepsilon(\delta)$

An IOPP for RS look like this:

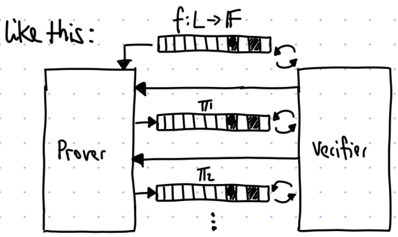

The efficiency measures are as in an IOP except we also charge for queries to f.

Henceforth we restrict our attention to smooth domains: $L = \langle \omega \rangle$ with $\mathrm{ord}(\omega) = 2^k$ as a subgroup of $\mathbb{F}^*$

<u>theorem</u>: For every $\mathbb{F}$, smooth domain $L \subseteq \mathbb{F}$ [⊛], and $d < |L|$,

$$RS[\mathbb{F},L,d] \in \text{IOPP}\left[\begin{array}{l} \varepsilon_c = 0,\ k = O(\log d),\ \ell = O(|L|),\ pt = O(|L|) \\ \varepsilon_s(\delta) = \text{"}1-\delta\text{"},\ q = O(\log d),\ vt = O(\log|L|),\ r = O(\log d) \end{array}\right]$$

this is called FRI protocol (Fast Reed-Solomon IOPP)

This IOPP for RS is important in practice and raises many elegant questions in coding theory.

[⊛ Similar statements hold for other types of (multiplicative or additive) subgroups L.]

Attempt 1: Recurse on Each Subproblem

$P((\mathbb{F},L,d),f)$	$f: L \to \mathbb{F}$	$V((\mathbb{F},L,d))$
Compute $\hat{g} := \mathrm{even}(\hat{f})$ and $\hat{h} := \mathrm{odd}(\hat{f})$ and set $g := \hat{g}\|_{L^2}$ and $h := \hat{h}\|_{L^2}$	$\xrightarrow{g,h: L^2 \to \mathbb{F}}$	sample $\mu \in L$ and check $f(\mu) \overset{?}{=} g(\mu^2) + \mu h(\mu^2)$
	recurse to test that $g \in RS[\mathbb{F},L^2,d/2]$, $h \in RS[\mathbb{F},L^2,d/2]$	

<u>Problem</u>: linear number of queries $\left(q(d) = 3 + 2q(d/2) = \Theta(d)\right)$

<u>Problem</u>: it's not even a test because distance decays in each recursion

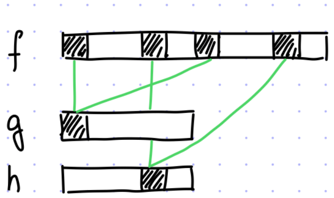

40 errors out of 100 locations $\delta = 40\%$
10 errors out of 50 locations $\delta = 20\%$
10 errors out of 50 locations $\delta = 20\%$

Such an example exists even if $\forall \mu \in L$ $f(\mu) = g(\mu^2) + \mu h(\mu^2)$!

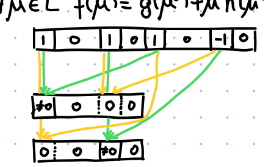

The distance could drop as $\delta \to \delta/2 \to \delta/4 \to \ldots \to \delta/2^r$.
We cannot sustain $r = \omega(1)$ rounds of interaction.

Attempt 2: Fold and Recurse [2/2]

$P((\mathbb{F},L,d),f)$	$f: L \to \mathbb{F}$	$V((\mathbb{F},L,d))$
Compute $\hat{g} := \mathrm{even}(\hat{f})$ and $\hat{h} := \mathrm{odd}(\hat{f})$ and set $g := \hat{g}\|_{L^2}$ and $h := \hat{h}\|_{L^2}$	$\xrightarrow{g,h: L^2 \to \mathbb{F}}$	sample $\mu \in L$ and check $f(\mu) \overset{?}{=} g(\mu^2) + \mu h(\mu^2)$
	$\xleftarrow{\alpha}$	sample $\alpha \in \mathbb{F}$
Set $f_\alpha := g + \alpha h$	$\xrightarrow{f_\alpha: L^2 \to \mathbb{F}}$	check $f_\alpha(\mu^2) = g(\mu^2) + \alpha h(\mu^2)$
	recurse to test that $f_\alpha \in RS[\mathbb{F},L^2,d/2]$	

Now consider the noisy case:
suppose f is δ-far from $RS[\mathbb{F},L,d]$.

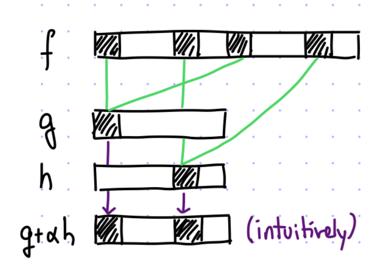

$\delta = 20\%$
$\delta = 10\%$
$\delta = 10\%$
$\delta = 20\%$

Folding seems to address the prior problem by preserving distance!

What if the cheating prover decreases distance by sending functions g, h, f_α that are inconsistent?

We do have consistency checks in each round for this.
So, informally, we have to (at least) pay an error of
$r \cdot \Pr[\text{a round's consistency check fails}]$
Since $r = \Theta(\log d)$ we have two options:
(i) make $\omega(1)$ queries/round (leads to $\omega(\log d)$ queries overall)
(ii) change the protocol

The FRI Protocol

Today we analyze the FRI protocol:

$P((\mathbb{F},L,d),f_0)$		$V((\mathbb{F},L,d))$
	$f_0 : L \to \mathbb{F}$	• interaction randomness: $\alpha_0,\alpha_1,\dots,\alpha_{r-1} \leftarrow \mathbb{F}$
	$\xleftarrow{\alpha_0}$	• consistency check randomness: $\mu \leftarrow L$ [can repeat t times]
$f_1 := \mathrm{Fold}(f_0,\alpha_0)$	$\xrightarrow{f_1 : L^2 \to \mathbb{F}}$	• consistency check on $f_0, f_1, \dots, f_r$:
	$\xleftarrow{\alpha_1}$	$f_1(\mu^2) \stackrel{?}{=} \frac{f_0(\mu)+f_0(-\mu)}{2} + \alpha_0 \cdot \frac{f_0(\mu)-f_0(-\mu)}{2\mu}$
$f_2 := \mathrm{Fold}(f_1,\alpha_1)$	$\xrightarrow{f_2 : L^4 \to \mathbb{F}}$	$f_2(\mu^4) \stackrel{?}{=} \frac{f_1(\mu^2)+f_1(-\mu^2)}{2} + \alpha_1 \cdot \frac{f_1(\mu^2)-f_1(-\mu^2)}{2\mu^2}$
	$\vdots$	$\vdots$
	$\xleftarrow{\alpha_{r-1}}$	$f_r(\mu^{2^r}) \stackrel{?}{=} \frac{f_{r-1}(\mu^{2^{r-1}})+f_{r-1}(-\mu^{2^{r-1}})}{2} + \alpha_{r-1} \cdot \frac{f_{r-1}(\mu^{2^{r-1}})-f_{r-1}(-\mu^{2^{r-1}})}{2\mu^{2^{r-1}}}$
$f_r := \mathrm{Fold}(f_{r-1},\alpha_{r-1})$	$\xrightarrow{f_r : L^{2^r} \to \mathbb{F}}$	• last function is low degree: $f_r \stackrel{?}{\in} RS[\mathbb{F}, L^{2^r}, d/2^r]$

query pattern:

$f_0(\mu)$ $f_0(-\mu)$

$f_1(\mu^2)$ $f_1(-\mu^2)$

$f_2(\mu^4)$ $f_2(-\mu^4)$

$\vdots$

$f_{r-1}(\mu^{2^{r-1}})$ $f_{r-1}(-\mu^{2^{r-1}})$

$f_r(\mu^{2^r})$

theorem: If $f_0 : L \to \mathbb{F}$ is δ-far from $RS[\mathbb{F},L,d]$ then $\forall \tilde{P}$

$$\Pr_{\alpha_0,\dots,\alpha_{r-1}}\left[\Pr_{\bar{\mu}\in L^t}\left[\langle \tilde{P}, V^f(\alpha,\bar{\mu})\rangle = 1\right] \le \left(1-\min\left\{\delta, \tfrac{1-\rho}{2}, \delta^*(\rho)\right\}\right)^t\right] \ge 1-\Omega\left(\frac{|L|}{|\mathbb{F}|}\right).$$

Here $\delta^*(\rho)$ is a universal constant with a dependence on the rate $\rho := d/|L|$.

In particular the soundness error is at most $O\left(\frac{|L|}{|\mathbb{F}|}\right) + \left(1-\min\left\{\delta, \frac{1-\rho}{2}, \delta^*(\rho)\right\}\right)^t$.

Soundness Analysis: Notations and Definitions

For notational simplicity: $L_i := L^{2^i}$, $d_i := d/2^i$, $\mu_i := \mu^{2^i}$.

Note that the rate is the same in each round's code: $\frac{d_i}{|L_i|} = \frac{d/2^i}{|L^{2^i}|} = \frac{d/2^i}{|L|/2^i} = \frac{d}{|L|} \triangleq \rho$

The (relative) distance between any two codewords in $RS[\mathbb{F}, L_i, d_i]$ is at least $1-\rho$.

Fix $f_0 : L \to \mathbb{F}$ and a prover $\tilde{P}$.

The prover $\tilde{P}$ is fully specified by functions $\{f_i : L_i \to \mathbb{F}\}_{i=1}^{r}$ with f_i depending on $\alpha_0, \dots, \alpha_{i-1} \in \mathbb{F}$.

Define $\forall i \in \{0,1,\dots,r-1\}$ $\mathrm{Fail}_i := \{a \in L_i \mid f_{i+1}(a^2) \neq \mathrm{Fold}(f_i, \alpha_i)(a)\}$.

Distance "by cosets": given $g, h : L_i \to \mathbb{F}$, $\Delta(g,h) := \frac{|\{a \in L_i \mid g(a) \neq h(a) \text{ or } g(-a) \neq h(-a)\}|}{|L_i|}$.

We keep track of distances for each round $i \in \{0,1,\dots,r\}$:

- $\delta_i \triangleq \Delta(f_i, RS[\mathbb{F}, L_i, d_i])$ fraction of cosets $\{-a,a\}$ to be changed for degree $< d_i$
- $\hat{f}_i$ is closest polynomial of degree $< d_i$ to $f_i : L_i \to \mathbb{F}$ (as measured by Δ)
- $\mathrm{Err}_i := \{a \in L_i \text{ s.t. } f_i(a) \neq \hat{f}_i(a) \text{ or } f_i(-a) \neq \hat{f}_i(-a)\}$.

If $\delta_i < \frac{1-\rho}{2}$ then $\hat{f}_i$ is unique and so Err_i is well-defined.

Soundness Analysis: Distortion

We have intuitively argued that random folding preserves distance with high probability.
Let's now formalize what we mean:

<u>def</u>: Given $f : L \to \mathbb{F}$ and $\delta \in (0,1)$ regular pointwise distance $\rho := d/|L|$

$$\mathrm{Drop}(f,\delta) := \{\alpha \in \mathbb{F} \mid \Delta(\mathrm{Fold}(f,\alpha), RS[\mathbb{F}, L^2, d/2]) < \delta\}.$$

<u>theorem</u>: Fix $f : L \to \mathbb{F}$ and set $\delta := \blacktriangle(f, RS[\mathbb{F}, L, d])$ (blockwise). Define $\delta^*(\rho) := \frac{1-5\rho}{4}$

① if $\delta < \frac{1-\rho}{2}$ then $\Pr_\alpha[\alpha \in \mathrm{Drop}(f,\delta)] \leq |L|/|\mathbb{F}|$

② if $\delta \geq \frac{1-\rho}{2}$ then $\Pr_\alpha[\alpha \in \mathrm{Drop}(f,\delta^*(\rho))] \leq |L|/|\mathbb{F}|$.

Hence, in the FRI protocol, the probability that some distortion happens is:

$$\Pr_{\alpha_0,\dots,\alpha_{r-1}}\left[\exists i \in \{0,1,\dots,r-1\} : \alpha_i \in \mathrm{Drop}(f_i, \min\{\delta_i, \delta^*(\rho)\})\right] \leq \sum_{i=0}^{r-1} \frac{|L_i|}{|\mathbb{F}|} = \left(\sum_{i=0}^{r-1} \frac{1}{2^i}\right)\frac{|L|}{|\mathbb{F}|} \leq \frac{2|L|}{|\mathbb{F}|}.$$

We take a union bound on this bad event, and henceforth assume that no distortion happens.
We wish to prove that $\Pr_\mu[\text{reject}] \geq \min\{\delta_0, \text{constants}\}$ when $\alpha_0,\dots,\alpha_{r-1}$ gives no distortion.

Soundness Analysis: Easy Case [1/2]

Suppose that $\tilde{P}$ adopts a "consistent but noisy" strategy.

That is, the interaction randomness $\alpha_0, \alpha_1, \ldots, \alpha_{r-1} \in \mathbb{F}$ is such that

① all functions are within unique decoding $\delta_0, \delta_1, \ldots, \delta_{r-1} < \frac{1-\rho}{2}$ ($\delta_r = 0$ always)

AND

② the (unique) corrections are consistent $\mathrm{Fold}(\hat{f}_0, \alpha_0) \equiv \hat{f}_1, \ldots, \mathrm{Fold}(\hat{f}_{r-1}, \alpha_{r-1}) \equiv \hat{f}_r$

lemma: $\Pr_{\mu}[\text{reject}] \geq \frac{|\mathrm{Err}_0|}{|L|} = \delta_0$

Recall: $\mathrm{Err}_i := \{a \in L_i \mid f_i(a) \neq \hat{f}_i(a) \text{ or } f_i(-a) \neq \hat{f}_i(-a)\}$

proof: Suppose WLOG that $\hat{f}_0$ is 0 on L_0. (If not, subtract $\hat{f}_0$ from f_0.)

By ②, we know that: $\hat{f}_1$ is 0 on L_1, $\hat{f}_2$ is 0 on $L_2, \ldots, \hat{f}_r$ is 0 on L_r.

Also, $f_r : L_r \to \mathbb{F}$ is 0 because $\delta_r = 0$ and so $f_r = \hat{f}_r|_{L_r} = 0$.

Fix $\mu_0 \in \mathrm{Err}_0 \subseteq L_0$ (which determines $\mu_1, \ldots, \mu_r$).

Let $j \in \{0, 1, \ldots, r\}$ be the largest index s.t. $\mu_j \in \mathrm{Err}_j \subseteq L_j$. (exists because $j=0$ is an option)

Note that $j < r$ because $f_r = \hat{f}_r|_{L_r}$ so that $\mathrm{Err}_r = \emptyset$.

By maximality of j, $\mu_{j+1} \notin \mathrm{Err}_{j+1}$ so $f_{j+1}(\mu_{j+1}) = \hat{f}_{j+1}(\mu_{j+1}) = 0$.

claim: $\mathrm{Fold}(f_j, \alpha_j)(\mu_{j+1}) \neq \mathrm{Fold}(\hat{f}_j, \alpha_j)(\mu_{j+1}) = 0$ [here we use $\alpha_j \notin \mathrm{Drop}(f_j, \delta_j)$, $\mu_j \in \mathrm{Err}_j$, & ①]

Hence $\mathrm{Fold}(f_j, \alpha_j)(\mu_{j+1}) \neq f_{j+1}(\mu_{j+1})$ so the verifier rejects. ∎

Soundness Analysis: Easy Case [2/2]

Suppose that $\tilde{P}$ adopts a "consistent but noisy" strategy.

That is, the interaction randomness $\alpha_0, \alpha_1, \dots, \alpha_{r-1} \in \mathbb{F}$ is such that

① all functions are within unique decoding $\delta_0, \delta_1, \dots, \delta_{r-1} < \frac{1-\rho}{2}$ ($\delta_r = 0$ always) AND ② the (unique) corrections are consistent $\mathrm{Fold}(\hat{f}_0, \alpha_0) \equiv \hat{f}_1, \dots, \mathrm{Fold}(\hat{f}_{r-1}, \alpha_{r-1}) \equiv \hat{f}_r$

claim: $\mathrm{Fold}(f_j, \alpha_j)(\mu_{j+1}) \neq \mathrm{Fold}(\hat{f}_j, \alpha_j)(\mu_{j+1}) = 0$ [here we use $\alpha_j \notin \mathrm{Drop}(f_j, \delta_j)$, $\mu_j \in \mathrm{Err}_j$, & ①]

proof:

- For every $a \notin \mathrm{Err}_j$, $\mathrm{Fold}(f_j, \alpha_j)(a^2) = \frac{f_j(a) + f_j(-a)}{2} + \alpha_j \frac{f_j(a) - f_j(-a)}{2a} = \frac{\hat{f}_j(a) + \hat{f}_j(-a)}{2} + \alpha_j \frac{\hat{f}_j(a) - \hat{f}_j(-a)}{2a} = \mathrm{Fold}(\hat{f}_j, \alpha_j)(a^2)$.
  Hence $\mathrm{Fold}(f_j, \alpha_j)$ and $\mathrm{Fold}(\hat{f}_j, \alpha_j)$ differ in at most $\frac{1}{2}|\mathrm{Err}_j| = \frac{1}{2}\delta_j |L_j| = \delta_j |L_{j+1}|$ locations on L_{j+1}.
  This implies that $\widehat{\mathrm{Fold}(f_j, \alpha_j)} \equiv \mathrm{Fold}(\hat{f}_j, \alpha_j)$ because they differ in at most $\delta_j |L_{j+1}| < \frac{1-\rho}{2} |L_{j+1}|$ locations.

- For every $a \in \mathrm{Err}_j$ (i.e., $f_j(a) \neq \hat{f}_j(a)$ or $f_j(-a) \neq \hat{f}_j(-a)$) if α_j is such that $\mathrm{Fold}(f_j, \alpha_j)(a^2) = \mathrm{Fold}(\hat{f}_j, \alpha_j)(a^2)$
  then $\Delta(\mathrm{Fold}(f_j, \alpha_j), RS[\mathbb{F}, L_j, d_j]) = \Delta(\mathrm{Fold}(f_j, \alpha_j), \widehat{\mathrm{Fold}(f_j, \alpha_j)}) = \Delta(\mathrm{Fold}(f_j, \alpha_j), \mathrm{Fold}(\hat{f}_j, \alpha_j)) < \delta_j$,
  which means that $\alpha_j \in \mathrm{Drop}(f_j, \delta_j)$ [α_j causes distortion].

- We have assumed that $\mu_j \in \mathrm{Err}_j$ and $\alpha_j \notin \mathrm{Drop}(f_j, \alpha_j)$ so we conclude that

$$\mathrm{Fold}(f_j, \alpha_j) \text{ and } \mathrm{Fold}(\hat{f}_j, \alpha_j) \text{ disagree at } \mu_j^2 = \mu_{j+1}. \quad \blacksquare$$

Soundness Analysis: Harder Case [1/2]

Suppose that $\tilde{P}$ jumps to "a far or inconsistent function".

That is, the interaction randomness $\alpha_0, \alpha_1, \ldots, \alpha_{r-1} \in \mathbb{F}$ is such that

① at least one function is far — $\exists i \in \{0,1,\ldots,r-1\}$ $\delta_i \geq \frac{1-\rho}{2}$ ($\delta_r = 0$ always)

OR

② the (unique) correction of a close function is inconsistent — $\exists i \in \{0,1,\ldots,r-1\}$ $\delta_i < \frac{1-\rho}{2}$ and $\mathrm{Fold}(\hat{f}_i, \alpha_i) \not\equiv \hat{f}_{i+1}$

<u>lemma:</u> $\Pr_{\mu}[\text{reject}] \geq \min\left\{ \frac{1-\rho}{2}, \delta^*(\rho) \right\}$

Recall: $\mathrm{Err}_i := \{ a \in L_i \mid f_i(a) \neq \hat{f}_i(a) \text{ or } f_i(-a) \neq \hat{f}_i(-a) \}$

$\mathrm{Fail}_i := \{ a \in L_i \mid f_{i+1}(a^2) \neq \mathrm{Fold}(f_i, \alpha_i)(a) \}$

<u>proof:</u> Let i be the largest index for which the above holds.

This means that $\delta_{i+1} < \frac{1-\rho}{2}$ so $\hat{f}_{i+1}$ and Err_{i+1} are well-defined.

<u>claim:</u> $\frac{|\mathrm{Fail}_{i+1} \cup \mathrm{Err}_{i+1}|}{|L_{i+1}|} \geq \min\left\{ \frac{1-\rho}{2}, \delta^*(\rho) \right\}$ [proved in next slide]

Fix any $\mu_0 \in L_0$, which induces $\mu_1, \mu_2, \ldots, \mu_r$.

- If $i+1 = r$ then $\mathrm{Err}_{i+1} = \emptyset$ so "$\mu_{i+1} \in \mathrm{Fail}_{i+1} \cup \mathrm{Err}_{i+1}$" implies that $\mu_{i+1} \in \mathrm{Fail}_{i+1}$ and so the verifier rejects.
- If $i+1 < r$ then $\alpha_{i+1}, \ldots, \alpha_{r-1}$ are such that:

① $\delta_{i+1}, \ldots, \delta_{r-1} < \frac{1-\rho}{2}$ AND ② $\mathrm{Fold}(\hat{f}_{i+1}, \alpha_{i+1}) \equiv \hat{f}_{i+2}, \ldots, \mathrm{Fold}(\hat{f}_{r-1}, \alpha_{r-1}) \equiv \hat{f}_r$

If $\mu_{i+1} \in \mathrm{Err}_{i+1}$ then similarly to the easy case we can conclude that the verifier rejects.

If $\mu_{i+1} \in \mathrm{Fail}_{i+1}$ then (trivially) the verifier rejects. Either way, "$\mu_{i+1} \in \mathrm{Fail}_{i+1} \cup \mathrm{Err}_{i+1}$" $\Rightarrow$ verifier rejects ∎

Soundness Analysis: Harder Case

Suppose that $\tilde{P}$ jumps to "a far or inconsistent function".

That is, the interaction randomness $\alpha_0, \alpha_1, \dots, \alpha_{r-1} \in \mathbb{F}$ is such that

① at least one function is far $\exists i \in \{0,1,\dots,r-1\}\ \delta_i \geq \frac{1-\rho}{2}$ ($\delta_r = 0$ always)

OR

② the (unique) correction of a close function is inconsistent $\exists i \in \{0,1,\dots,r-1\}\ \delta_i < \frac{1-\rho}{2}$ and $\mathrm{Fold}(\hat{f}_i, \alpha_i) \neq \hat{f}_{i+1}$

claim: $\frac{|\mathrm{Fail}_{i+1} \cup \mathrm{Err}_{i+1}|}{|L_{i+1}|} \underset{ⓐ}{\geq} \Delta(\hat{f}_{i+1}|_{L_{i+1}}, \mathrm{Fold}(f_i, \alpha_i)) \underset{ⓑ}{\geq} \min\left\{\frac{1-\rho}{2}, \delta^*(\rho)\right\}$

Recall: $\mathrm{Err}_i := \{a \in L_i \mid f_i(a) \neq \hat{f}_i(a) \text{ or } f_i(-a) \neq \hat{f}_i(-a)\}$
$\mathrm{Fail}_i := \{a \in L_i \mid f_{i+1}(a^2) \neq \mathrm{Fold}(f_i, \alpha_i)(a)\}$

proof:

ⓐ If $\mu_{i+1} \in L_{i+1}$ is not in Err_{i+1} then $\hat{f}_{i+1}(\mu_{i+1}) = f_{i+1}(\mu_{i+1})$.
If $\mu_{i+1} \in L_{i+1}$ is not in Fail_{i+1} then $f_{i+1}(\mu_{i+1}) = \mathrm{Fold}(f_i, \alpha_i)(\mu_{i+1})$.

ⓑ If $\delta_i \geq \frac{1-\rho}{2}$ then (due to no distortion) $\mathrm{Fold}(f_i, \alpha_i)$ is $\delta^*(\rho)$-far from $RS[\mathbb{F}, L_{i+1}, d_{i+1}] \ni \hat{f}_{i+1}|_{L_{i+1}}$.
If $\delta_i < \frac{1-\rho}{2}$ then $\mathrm{Fold}(\hat{f}_i, \alpha_i) \neq \hat{f}_{i+1}$ so they differ in at least $\frac{|L_{i+1}| - d/2^{i+1}}{|L_{i+1}|} = 1-\rho$ locations.
Hence

$$1-\rho \leq \Delta(\hat{f}_{i+1}|_{L_{i+1}}, \mathrm{Fold}(\hat{f}_i, \alpha_i)|_{L_{i+1}}) \leq \Delta(\hat{f}_{i+1}|_{L_{i+1}}, \mathrm{Fold}(f_i, \alpha_i)) + \Delta(\mathrm{Fold}(f_i, \alpha_i), \mathrm{Fold}(\hat{f}_i, \alpha_i)|_{L_{i+1}})$$
$$= \Delta(\hat{f}_{i+1}|_{L_{i+1}}, \mathrm{Fold}(f_i, \alpha_i)) + \delta_i < \Delta(\hat{f}_{i+1}|_{L_{i+1}}, \mathrm{Fold}(f_i, \alpha_i)) + \frac{1-\rho}{2}.$$

We conclude that $\Delta(\hat{f}_{i+1}|_{L_{i+1}}, \mathrm{Fold}(f_i, \alpha_i)) \geq (1-\rho) - (\frac{1-\rho}{2}) = \frac{1-\rho}{2}$. ■

On Distortion for FRI

Fix $f: L \to \mathbb{F}$ and set $\delta := \Delta(f, RS[\mathbb{F}, L, d])$. Say that we want to prove that:

$$\Pr_{\alpha}[\alpha \in \mathrm{Drop}(f, \delta^*)] = \Pr_{\alpha}[\Delta(\mathrm{Fold}(f,\alpha), RS[\mathbb{F}, L^2, d/2]) < \delta^*] \leq \varepsilon$$

for desired δ^* and ε (that can be functions of $\delta, \mathbb{F}, \ldots$).

For this it suffices to prove statements such as the following:

Given a set $S \subseteq \mathbb{F}^n$, we write $S^{[m]}$ for the set of all matrices in $\mathbb{F}^{m\times n}$ whose rows are in S.

Then for $V = \begin{pmatrix} - v_1 - \\ \vdots \\ - v_m - \end{pmatrix} \in \mathbb{F}^{m \times n}$, $\Delta(V, S^{[m]}) :=$ "min fraction of cols in V to change to get elt in $S^{[m]}$".

<u>template lemma</u>: Fix $v_1, \ldots, v_m \in \mathbb{F}^n$ and a subspace $S \subseteq \mathbb{F}^n$ s.t. $\Delta(V, S^{[m]}) \geq \delta$

Then $\Pr_{\alpha_1,\ldots,\alpha_m}[\Delta(\alpha_1 v_1 + \cdots + \alpha_m v_m, S) < \delta^*] \leq \varepsilon$.

The goal follows by setting $S := RS[\mathbb{F}, L^2, d/2]$, $v_1(a^2) := \frac{f(a)+f(-a)}{2}$, $v_2(a^2) := \frac{f(a)-f(-a)}{2a}$:

① $\Delta(\alpha_1 v_1 + \alpha_2 v_2, S) = \Delta(v_1 + \frac{\alpha_2}{\alpha_1} v_2, S)$ $\forall (\alpha_1, \alpha_2) \in \mathbb{F}^2$ with $\alpha_1 \neq 0$

② $\Delta(f, RS[\mathbb{F}, L, d]) \geq \delta \to \Delta\left(\begin{bmatrix} - v_1 - \\ - v_2 - \end{bmatrix}, S^{[2]}\right) \geq \delta$

[if $\begin{bmatrix} - v_1 - \\ - v_2 - \end{bmatrix}$ differs in $<\delta$ columns with $\begin{bmatrix} - \hat{v}_1 - \\ - \hat{v}_2 - \end{bmatrix} \in S^{[2]}$ then $\hat{f}(x) := \hat{v}_1(x^2) + x\hat{v}_2(x^2)$ has $\deg < d$ and differs in $<\delta$ cosets of L with f]

[the probability goes from ε to $\frac{|\mathbb{F}|}{|\mathbb{F}|-1} \cdot \varepsilon$]

Course outline

Local-to-global phenomena

- Linearity testing
- Low-degree testing
- FFT-based testing of univariate polynomials

PCP constuctions

- exp-size, $O(1)$-local PCPs
- poly-size, polylog-local PCPs
- PCP composition
- Sublinear-time verification

Applications

- Delegation of computation
- Hardness of approximation