Foundations of Probabilistic Proofs

A course by **Alessandro Chiesa**

Lecture 16

Analysis of FRI

The FRI Protocol

Recall the FRI protocol to test proximity to $RS[\mathbb{F},L,d]=\{f:L\to\mathbb{F} : \deg(\hat{f})<d\}$:

(L: subgroup of $\mathbb{F}^*$ whose size is a power of 2)

$P((\mathbb{F},L,d),f_0)$		$V((\mathbb{F},L,d))$
	$f_0:L\to\mathbb{F}$	
	$\xleftarrow{\alpha_0\in\mathbb{F}}$	Interaction randomness: $\alpha_0,\alpha_1,\dots,\alpha_{r-1}\leftarrow\mathbb{F}$
$f_1:=\text{Fold}(f_0,\alpha_0)$	$\xrightarrow{f_1:L^2\to\mathbb{F}}$	Consistency check randomness: $\mu_1,\dots,\mu_t\leftarrow L$
	$\xleftarrow{\alpha_1\in\mathbb{F}}$	For each repetition $j\in[t]$:
$f_2:=\text{Fold}(f_1,\alpha_1)$	$\xrightarrow{f_2:L^4\to\mathbb{F}}$	$\forall\, i\in\{0,1,\dots,r-1\}$
	$\vdots$	$f_{i+1}(\mu_j^{2^{i+1}}) \overset{?}{=} \frac{f_i(\mu_j^{2^i})+f_i(-\mu_j^{2^i})}{2}+\alpha_i\cdot\frac{f_i(\mu_j^{2^i})-f_i(-\mu_j^{2^i})}{2\mu_j^{2^i}}$
	$\xleftarrow{\alpha_{r-1}\in\mathbb{F}}$	
$f_r:=\text{Fold}(f_{r-1},\alpha_{r-1})$	$\xrightarrow{f_r:L^{2^r}\to\mathbb{F}}$	Low-degree check: $f_r \overset{?}{\in} RS[\mathbb{F},L^{2^r},d/2^r]$.

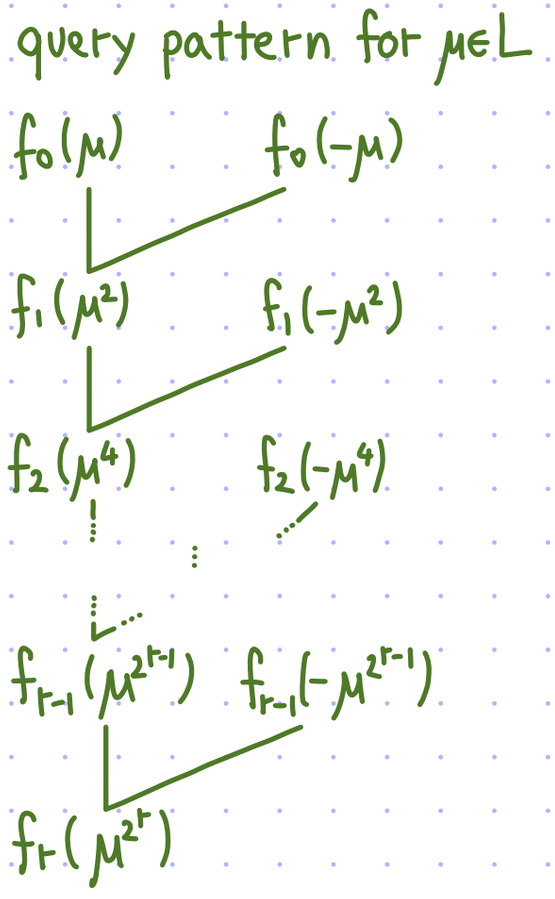

We prove that the soundness error is $\leq O\left(\frac{|L|}{|\mathbb{F}|}\right)+\left(1-\min\{\delta,c(\frac{d}{|L|})\}\right)^t$.

theorem: If $f_0:L\to\mathbb{F}$ is δ-far from $RS[\mathbb{F},L,d]=\{f:L\to\mathbb{F} : \deg(\hat{f})<d\}$ then

$$\forall\,\tilde{P}\quad \Pr_{\alpha_0,\dots,\alpha_{r-1}\in\mathbb{F}}\left[\Pr_{\mu_1,\dots,\mu_t\in L}\left[\langle\tilde{P},V^f(\vec{\alpha},\vec{\mu})\rangle=1\right]\leq\left(1-\min\{\delta,c(\tfrac{d}{|L|})\}\right)^t\right]\geq 1-\frac{2|L|}{|\mathbb{F}|}.$$

$c(\frac{d}{|L|})$: constant that depends on the rate $\rho:=\frac{d}{|L|}$

Soundness Analysis: Notations

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

The rate is the same in each round's RS code:

$$\frac{d_i}{|L_i|} = \frac{d/2^i}{|L^{2^i}|} = \frac{d/2^i}{|L|/2^i} = \frac{d}{|L|} = \rho$$

The (relative) distance between any two codewords in $RS[\mathbb{F}, L_i, d_i]$ is $\geq 1 - \frac{d_i - 1}{|L_i|} = 1 - \rho + \frac{1}{|L_i|} \geq 1 - \rho$.

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

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

Distance "by cosets": given $g, h : L_i \to \mathbb{F}$,

$$\Delta(g,h) := \frac{|\{a \in L_i \mid g(a) \neq h(a) \text{ OR } g(-a) \neq h(-a)\}|}{|L_i|}.$$ [Note that $\Delta(g,h) \geq \Delta(g,h)$.]

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

$\delta_i := \Delta(f_i, RS[\mathbb{F}, L_i, d_i])$ ← fraction of cosets $\{-a, a\}$ to change for degree $< d_i$

We denote by $\hat{f}_i$ a polynomial of degree $< d_i$ closest to $f_i : L_i \to \mathbb{F}$ (as measured by Δ).

Soundness Analysis: Distortion

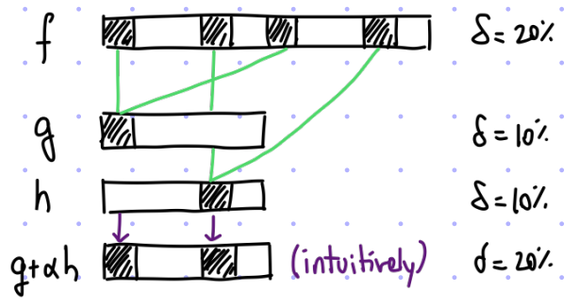

We saw intuition for how random folding preserves distance whp:
We formalize this intuition.

<u>def</u>: For $f: L \to \mathbb{F}$ and $\delta \in (0,1)$

(regular) pointwise distance

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

<u>lemma</u>: Fix $f: L \to \mathbb{F}$ and set $\delta := \Delta^{\text{(blockwise)}}(f, RS[\mathbb{F}, L, d])$.

$\rho := \frac{d}{|L|}$ (rate)

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

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

$\delta^*(\rho) := \frac{1-\rho}{8}$

The probability that distortion happens in the FRI protocol is:

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

We take a union bound on this, and henceforth assume that no distortion happens.

When $\alpha_0,\ldots,\alpha_{r-1}$ give no distortion, we wish to prove that $\Pr_M[\text{reject}] \ge \min\{\delta_0, \text{constants}\}$.

Soundness Analysis: Two Cases

Given that no distortion happens,

how can a prover attack the FRI protocol?

Key definitions for analysis:

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

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

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

(for $i \in \{0, 1, \ldots, r-1\}$)

① The "consistent but noisy" strategy.

Informally, $f_1, \ldots, f_r$ are close to RS and such that $\text{Fold}(\hat{f}_0, \alpha_0) \equiv \hat{f}_1, \ldots, \text{Fold}(\hat{f}_{r-1}, \alpha_{r-1}) \equiv \hat{f}_r$.

We prove that $\Pr_M[\text{reject}] \geq \frac{|\text{Noise}_0|}{|L_0|} \geq \delta_0$. (Note that $\frac{|\text{Noise}_i|}{|L_i|} = \Delta(f_i, \hat{f}_i) \geq \Delta(f_i, RS[\mathbb{F}, L_i, d_i]) = \delta_i$.)

② The "far or inconsistent" strategy.

Informally, $\exists$ (maximal) $i \in \{0, 1, \ldots, r-1\}$ s.t. f_i is far from RS or $\text{Fold}(\hat{f}_i, \alpha_i) \neq \hat{f}_{i+1}$.

We prove that $\Pr_M[\text{reject}] \geq \frac{|\text{Noise}_{i+1} \cup \text{Fail}_{i+1}|}{|L_{i+1}|} \geq \min\{\frac{1-\rho}{2}, \delta^*(\rho)\}$.

Soundness Analysis: Case 1 [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 AND ② the (unique) corrections are consistent

$\delta_0, \delta_1, \ldots, \delta_{r-1} < \frac{1-\rho}{2}$ ($\delta_r = 0$ always)

$\mathrm{Fold}(\hat{f}_0, \alpha_0) \equiv \hat{f}_1, \ldots, \mathrm{Fold}(\hat{f}_{r-1}, \alpha_{r-1}) \equiv \hat{f}_r$

lemma: $\Pr_{V}[\text{reject}] \geq \frac{|\mathrm{Noise}_0|}{|L_0|} \geq \delta_0$ (Recall: $\mathrm{Noise}_i := \{a \in L_i \mid f_i(a) \neq \hat{f}_i(a) \text{ or } f_i(-a) \neq \hat{f}_i(-a)\}$.)

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

By ②, $\hat{f}_1$ is 0 on L_1, $\hat{f}_2$ is 0 on L_2, …, $\hat{f}_r$ 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 $\mu_0 \in \mathrm{Noise}_0 \subseteq L_0$ (which determines $\mu_1, \ldots, \mu_r$).

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

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

By maximality of j, $\mu_{j+1} \notin \mathrm{Noise}_{j+1}$ 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 $\alpha_j \notin \mathrm{Drop}(f_j, \delta_j)$, $\mu_j \in \mathrm{Noise}_j$, & ①]

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

Soundness Analysis: Case 1 [2/2]

Suppose that $\widetilde{P}$ 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 AND ② the (unique) corrections are consistent

$\delta_0, \delta_1, \ldots, \delta_{r-1} < \frac{1-\rho}{2}$ ($\delta_r = 0$ always)

$\mathrm{Fold}(\hat{f}_0, \alpha_0) \equiv \hat{f}_1, \ldots, \mathrm{Fold}(\hat{f}_{r-1}, \alpha_{r-1}) \equiv \hat{f}_r$

Left to prove: claim: $\mathrm{Fold}(f_j, \alpha_j)(\mu_{j+1}) + \mathrm{Fold}(\hat{f}_j, \alpha_j)(\mu_{j+1}) = 0$

[Recall: $\mathrm{Noise}_i = \{a \in L_i \mid f_i(a) \neq \hat{f}_i(a) \text{ OR } f_i(-a) \neq \hat{f}_i(-a)\}$]

proof:

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

- For every $a \in \mathrm{Noise}_j$ (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 $\Delta(\mathrm{Fold}(f_j, \alpha_j), \mathrm{RS}[\mathbb{F}, L_j, d_j]) = \Delta(\mathrm{Fold}(f_j, \alpha_j), \widetilde{\mathrm{Fold}(f_j, \alpha_j)}) = \Delta(\mathrm{Fold}(f_j, \alpha_j), \mathrm{Fold}(\hat{f}_j, \alpha_j)) < \delta_j$,
  which means that $\alpha_j \in \mathrm{Drop}(f_j, \delta_j)$ [α_j causes distortion].

- We assumed that $\mu_j \in \mathrm{Noise}_j$ and $\alpha_j \notin \mathrm{Drop}(f_j, \alpha_j)$ so we conclude that
  $\mathrm{Fold}(f_j, \alpha_j)$ and $\mathrm{Fold}(\hat{f}_j, \alpha_j)$ disagree at $\mu_j^2 = \mu_{j+1}$. ■

Soundness Analysis: Case 2 [1/2]

Suppose that $\widetilde{P}$ 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 — OR — ② the (unique) correction of a close function is inconsistent

$\exists i \in \{0,1,\ldots,r-1\}\ \delta_i \geq \frac{1-\rho}{2}$ $(\delta_r = 0$ always$)$

$\exists i \in \{0,1,\ldots,r-1\}\ \delta_i < \frac{1-\rho}{2}$ and $\mathrm{Fold}(\hat{f}_i, \alpha_i) \neq \hat{f}_{i+1}$

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

Recall: $\mathrm{Noise}_i := \{a \in L_i \mid f_i(a) \neq \hat{f}_i(a)$ 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: 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 Noise_{i+1} are well-defined.

claim: $\frac{|\mathrm{Fail}_{i+1} \cup \mathrm{Noise}_{i+1}|}{|L_{i+1}|} \geq \min\left\{\frac{1-\rho}{2}, \delta^*(\rho)\right\}$ [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 $\mathrm{Noise}_{i+1} = \emptyset$ so $\mu_{i+1} \in \mathrm{Fail}_{i+1} \cup \mathrm{Noise}_{i+1}$ 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 $\mathrm{Fold}(\hat{f}_{i+1}, \alpha_{i+1}) = \hat{f}_{i+2}, \ldots, \mathrm{Fold}(\hat{f}_{r-1}, \alpha_{r-1}) = \hat{f}_r$

If $\mu_{i+1} \in \mathrm{Noise}_{i+1}$ then similarly to Case 1 we conclude that the verifier rejects.

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

Soundness Analysis: Case 2 [2/2]

Suppose that $\widetilde{P}$ 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 — OR — ② the (unique) correction of a close function is inconsistent

① $\exists i \in \{0,1,\ldots,r-1\}\ \delta_i \geq \frac{1-\rho}{2}$ ($\delta_r = 0$ always)

② $\exists i \in \{0,1,\ldots,r-1\}\ \delta_i < \frac{1-\rho}{2}$ and $\mathrm{Fold}(\hat{f}_i, \alpha_i) \neq \hat{f}_{i+1}$

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

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

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

proof:

@ If $\mu_{i+1} \in L_{i+1}$ is not in Noise_{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 $\mu_{i+1} \notin \mathrm{Fail}_{i+1} \cup \mathrm{Noise}_{i+1}$ then $\hat{f}_{i+1}(\mu_{i+1}) = \mathrm{Fold}(f_i, \alpha_i)(\mu_{i+1})$.

ⓑ 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 agree in $< d/2^{i+1} = \rho \cdot |L_{i+1}|$ 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

We return to the discussion of DISTORTION.

(regular) pointwise distance

<u>def</u>: For $f: L \to \mathbb{F}$ and $\delta \in (0,1)$, $\mathrm{Drop}(f,\delta) := \{\alpha \in \mathbb{F} \mid \Delta(\mathrm{Fold}(f,\alpha), \mathrm{RS}[\mathbb{F}, L^2, \frac{d}{2}]) < \delta\}$.

We wish to prove statements such as:

blockwise

Let $f: L \to \mathbb{F}$ and $\delta := \Delta(f, \mathrm{RS}[\mathbb{F}, L, d])$. Then $\Pr_{\alpha}[\alpha \in \mathrm{Drop}(f, \delta^*)] \leq \epsilon$.

Here $\delta^* \in (0, \delta]$ and $\epsilon \in (0,1)$ can be functions of $\delta, \mathbb{F}, \ldots$

This is proved via two statements in different regimes.

FRI Distortion Lemma

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

blockwise

① small distance case: $\delta < \frac{1-\rho}{2} \to \Pr_{\alpha \leftarrow \mathbb{F}}[\alpha \in \mathrm{Drop}(f, \delta)] \leq \frac{|L|}{|\mathbb{F}|}$.

② large distance case: $\delta \geq \frac{1-\rho}{2} \to \Pr_{\alpha \leftarrow \mathbb{F}}[\alpha \in \mathrm{Drop}(f, \frac{1-\rho}{8})] \leq \frac{1}{|\mathbb{F}|}$.

We prove ② and then ①.

A Template Statement

Distortion is related to WORST-VS-AVERAGE-CASE DISTANCES TO SUBSPACES.

<u>def:</u>
- m-wise interleaving of a set $S \subseteq \mathbb{F}^n$: $S^{[m]} := \{ A \in \mathbb{F}^{m \times n} \mid \text{each row of } A \text{ is in } S \}$.
- column distance: $\Delta_{col}(A \in \mathbb{F}^{m \times n}, B \in \mathbb{F}^{m \times n}) :=$ "fraction of columns at which A and B differ"

<u>Template Statement:</u> Let $S \subseteq \mathbb{F}^n$ be a subspace and $v_1, \dots, v_m \in \mathbb{F}^n$ be vectors.
If $V := \begin{bmatrix} — & v_1 & — \\ & \vdots & \\ — & v_m & — \end{bmatrix} \in \mathbb{F}^{m \times n}$ is such that $\Delta_{col}(V, S^{[m]}) \geq \delta$ then $\Pr_{\alpha_1, \dots, \alpha_m \leftarrow \mathbb{F}} [\Delta(\alpha_1 v_1 + \dots + \alpha_m v_m, S) < \delta^*] \leq \varepsilon$.

The template suffices to establish FRI distortion with the same δ^* and ε:

- define: $S := RS[\mathbb{F}, L^2, \frac{d}{2}]$, $v_1 := \left(\frac{f(a)+f(-a)}{2}\right)_{a^2 \in L^2}$, $v_2 := \left(\frac{f(a)-f(-a)}{2a}\right)_{a^2 \in L^2}$
- <u>claim:</u> $\Delta(f, RS[\mathbb{F}, L, d]) \leq \Delta_{col}(\begin{bmatrix} — & v_1 & — \\ — & v_2 & — \end{bmatrix}, S^{[2]})$

  Let $\begin{bmatrix} — & \hat{v}_1 & — \\ — & \hat{v}_2 & — \end{bmatrix} \in S^{[2]}$ be such that $\Delta_{col}(\begin{bmatrix} — & v_1 & — \\ — & v_2 & — \end{bmatrix}, S^{[2]}) = \Delta_{col}(\begin{bmatrix} — & v_1 & — \\ — & v_2 & — \end{bmatrix}, \begin{bmatrix} — & \hat{v}_1 & — \\ — & \hat{v}_2 & — \end{bmatrix})$ for some $\hat{v}_1, \hat{v}_2 \in S$.

  Define $\hat{f} : L \to \mathbb{F}$ as $\hat{f}(a) := \hat{v}_1(a^2) + a\hat{v}_2(a^2)$. Note that $\hat{f} \in RS[\mathbb{F}, L, d]$ (so $\Delta(f, RS[\mathbb{F}, L, d]) \leq \Delta(f, \hat{f})$)

  and that $\forall a \in L$ if $v_1(a^2) = \hat{v}_1(a^2)$ and $v_2(a^2) = \hat{v}_2(a^2)$ then $f(a) = \hat{f}(a)$ (so $\Delta(f, \hat{f}) \leq \Delta_{col}(\begin{bmatrix} — & v_1 & — \\ — & v_2 & — \end{bmatrix}, \begin{bmatrix} — & \hat{v}_1 & — \\ — & \hat{v}_2 & — \end{bmatrix})$).

- $\Pr_{\alpha \leftarrow \mathbb{F}}[\alpha \in Drop(f, \delta^*)] = \Pr_{\alpha \leftarrow \mathbb{F}}[\Delta(v_1 + \alpha \cdot v_2, S) < \delta^*] = \Pr_{\alpha_1, \alpha_2 \leftarrow \mathbb{F}}[\Delta(\alpha_1 v_1 + \alpha_2 v_2, S) < \delta^* \wedge \alpha_1 \neq 0] \leq \Pr_{\alpha_1, \alpha_2 \leftarrow \mathbb{F}}[\Delta(\alpha_1 v_1 + \alpha_2 v_2, S) < \delta^*]$

Distortion with Half Distance

We start with a simpler statement:

lemma: Let $S \subseteq \mathbb{F}^n$ be a subspace and $v_1, \ldots, v_m \in \mathbb{F}^n$ be vectors.
If $\exists i \in [m]$ s.t. $\Delta(v_i, S) \geq \delta$ then $\Pr_{\alpha_1, \ldots, \alpha_m \leftarrow \mathbb{F}} [\Delta(\alpha_1 v_1 + \cdots + \alpha_m v_m, S) < \delta/2] \leq \frac{1}{|\mathbb{F}|}$.

Achieves template lemma with $\delta^* = \delta/2$ and $\epsilon = \frac{1}{|\mathbb{F}|}$ under a stronger assumption (implies that $\Delta(\bar{V}, S^{[m]}) \geq \delta$)

proof: Without loss of generality take $i=1$.
Fix arbitrary $\alpha_2, \ldots, \alpha_m \in \mathbb{F}$ and define $y := \alpha_2 v_2 + \cdots + \alpha_m v_m$.
Suppose by way of contradiction that $\exists$ distinct $\alpha_1, \alpha_1' \in \mathbb{F}$ and $\exists$ $w, w' \in S$ s.t.

$$\Delta(\alpha_1 v_1 + y, w) < \delta/2 \quad \text{and} \quad \Delta(\alpha_1' v_1 + y, w') < \delta/2.$$

Then
$$\begin{aligned}
\Delta(v_1, S) = \Delta((\alpha_1 - \alpha_1') \cdot v_1, S) &\leq \Delta((\alpha_1 - \alpha_1') \cdot v_1, w - w') \\
&= \Delta((\alpha_1 \cdot v_1 + y) - (\alpha_1' \cdot v_1 + y), w - w') \\
&\leq \Delta(\alpha_1 \cdot v_1 + y, w) + \Delta(\alpha_1' \cdot v_1 + y, w') \\
&< \delta/2 + \delta/2 = \delta.
\end{aligned}$$

We conclude that $\forall \alpha_2, \ldots, \alpha_m \in \mathbb{F}$ $\Pr_{\alpha_1} [\Delta(\alpha_1 v_1 + \cdots + \alpha_m v_m, S) < \delta/2] \leq \frac{1}{|\mathbb{F}|}$. ■

FRI Distortion: Large Distance Case

FRI Distortion Lemma — large distance case

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

If $\delta \geq \frac{1-\rho}{2}$ then $\Pr_{\alpha \leftarrow \mathbb{F}}\left[\alpha \in Drop\left(f, \frac{1-\rho}{8}\right)\right] \leq \frac{1}{|\mathbb{F}|}$.

Recall that $Drop(f,\delta)$ is $\{\alpha \in \mathbb{F} \mid \Delta(Fold(f,\alpha), RS[\mathbb{F}, L^2, \frac{d}{2}]) < \delta\}$ (regular) pointwise distance

proof:

We invoke the half distance lemma.

Define $S := RS[\mathbb{F}, L^2, \frac{d}{2}]$ $\quad v_1 := \left(\frac{f(a)+f(-a)}{2}\right)_{a^2 \in L^2}$ $\quad v_2 := \left(\frac{f(a)-f(-a)}{2a}\right)_{a^2 \in L^2}$

Note that $\Delta_{col}([\text{—}v_1\text{—}; \text{—}v_2\text{—}], S^{[2]}) \geq \Delta(f, RS[\mathbb{F}, L, d]) = \delta$ implies that $\exists i \in [2]$ s.t. $\Delta(v_i, S) \geq \delta/2$, because the fraction of column errors is at most the largest fraction of errors in a row times the number of rows.

Hence $\Pr_{\alpha \leftarrow \mathbb{F}}\left[\alpha \in Drop\left(f, \frac{1-\rho}{8}\right)\right] \leq \Pr_{\alpha \leftarrow \mathbb{F}}\left[\alpha \in Drop\left(f, \delta/4\right)\right]$ (because $\frac{1-\delta}{2} \leq \delta$, and $\gamma \leq \gamma' \to Drop(f,\gamma) \subseteq Drop(f,\gamma')$)

$$\leq \Pr_{\alpha_1, \alpha_2 \leftarrow \mathbb{F}}\left[\Delta(\alpha_1 v_1 + \alpha_2 v_2, S) < \frac{\delta}{4}\right] \leq \frac{1}{|\mathbb{F}|}.$$

proved 2 slides ago

distortion with half distance (view $\frac{\delta}{4}$ as $\frac{\delta/2}{2}$)

■

FRI Distortion: Small Distance Case

FRI Distortion Lemma — small distance case

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

If $\delta < \frac{1-\rho}{2}$ then $\Pr_{\alpha \leftarrow \mathbb{F}}[\alpha \in Drop(f, \delta)] \leq \frac{|L|}{2} \cdot \frac{1}{|\mathbb{F}|}$.

Recall that $Drop(f, \delta)$ is $\{\alpha \in \mathbb{F} \mid \Delta(Fold(f,\alpha), RS[\mathbb{F}, L^2, \frac{d}{2}]) < \delta\}$

(Δ: (regular) pointwise distance)

and $Fold(f,\alpha)(a^2) = \frac{f(a)+f(-a)}{2} + \alpha \cdot \frac{f(a)-f(-a)}{2a}$

proof: Let $\hat{f} \in RS[\mathbb{F}, L, d]$ be the unique closest codeword to f (wrt Δ). It suffices to show

$Drop(f,\delta) = \{\alpha \in \mathbb{F} \mid \exists a^2 \in Noise(f) : Fold(f,\alpha)(a^2) = Fold(\hat{f},\alpha)(a^2)\}$ where $Noise(f) := \left\{a^2 \in L^2 \mid \begin{matrix} f(a) \neq \hat{f}(a) \\ \vee\, f(-a) \neq \hat{f}(-a) \end{matrix}\right\}$.

Indeed:
- $\forall a^2 \in Noise(f)$, $Fold(f,\alpha)(a^2) - Fold(\hat{f},\alpha)(a^2)$ is a non-zero linear polynomial in α
- $\Pr_{\alpha \leftarrow \mathbb{F}}[\alpha \in Drop(f,\delta)] \leq \sum_{a^2 \in Noise(f)} \Pr_{\alpha \in \mathbb{F}}[Fold(f,\alpha)(a^2) = Fold(\hat{f},\alpha)(a^2)] \leq |Noise(f)| \cdot \frac{1}{|\mathbb{F}|} \leq \frac{|L|}{2} \cdot \frac{1}{|\mathbb{F}|}$

For every $a \in L$, if $f(a) = \hat{f}(a)$ and $f(-a) = \hat{f}(-a)$ then, $\forall \alpha \in \mathbb{F}$, $Fold(f,\alpha)(a^2) = Fold(\hat{f},\alpha)(a^2)$.

Hence $\Delta(f,\hat{f}) = \delta$ implies, $\forall \alpha \in \mathbb{F}$, that $\Delta(Fold(f,\alpha), Fold(\hat{f},\alpha)) \leq \delta$.

Since $\delta < \frac{1-\rho}{2}$, $\forall \alpha \in \mathbb{F}$, $Fold(\hat{f},\alpha) \in RS[\mathbb{F}, L^2, d/2]$ is the closest codeword to $Fold(f,\alpha)$ wrt Δ.

We conclude that $\alpha \notin Drop(f,\delta) \leftrightarrow \forall a^2 \in Noise(f), Fold(f,\alpha)(a^2) \neq Fold(\hat{f},\alpha)(a^2)$

Bibliography

FRI protocol

- [BBHR 2018]: Fast Reed–Solomon interactive oracle proofs of proximity, by Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, Michael Riabzev.
- [BKS 2018]: Worst-case to average case reductions for the distance to a code, by Eli Ben-Sasson, Swastik Kopparty, Shubhangi Saraf.
- [BGKS 2019]: DEEP-FRI: Sampling outside the box improves soundness, by Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty, Shubhangi Saraf. (▸Video)
- [BCIKS 2020]: Proximity gaps for Reed–Solomon codes, by Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty, Shubhangi Saraf. (▸Video 1), (▸Video 2)
- [ABN 2022]: Efficient multivariate low-degree tests via interactive oracle proofs of proximity for polynomial codes, by Daniel Augot, Sarah Bordage, Jade Nardi.
- [BLNR 2020]: Interactive oracle proofs of proximity to algebraic geometry codes, by Sarah Bordage, Mathieu Lhotel, Jade Nardi, Hughes Randriam. (▸Video)
- [ACFY 2024]: STIR: Reed–Solomon proximity testing with fewer queries, by Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, Eylon Yogev. (▸Video 1), (▸Video 2), (▸Video 3), (▸Podcast), (▸Blog)
- [ACFY 2024]: WHIR: Reed–Solomon proximity testing with super-fast verification, by Gal Arnon, Alessandro Chiesa, Giacomo Fenzi, Eylon Yogev. (▸Blog)

Extends soundness analysis of FRI to list-decoding regime.

Improves soundness with OOD and quotients

State-of-the-art analysis, removes need for quotients

Extends FRI to algebraic geometry codes

Recent IOPP for RS, with better query complexity

More recent IOPP for RS, with better query complexity and super fast verification