Foundations of Probabilistic Proofs

A course by **Alessandro Chiesa**

Lecture 15

FRI Protocol

Proximity Testing to the Reed–Solomon Code

We seek a proximity test V for $RS[\mathbb{F},L,d] := \{ f: L \to \mathbb{F} \text{ s.t. } \deg(\hat{f}) < d \}$:

① completeness: $\forall f \in RS[\mathbb{F},L,d] \quad \Pr[V^f = 1] = 1$

② soundness: $\forall f: L \to \mathbb{F}$, if f is δ-far from $RS[\mathbb{F},L,d]$ then $\Pr[V^f = 1] \leq \varepsilon(\delta)$ (s.t. $\delta = \Omega(1) \to \varepsilon(\delta) = O(1)$)

We have seen that:

- $d+1$ queries suffice to achieve $\varepsilon(\delta) = 1 - \delta$

  [interpolate the answers to any d queries, and check consistency with the answer to a random query]

- $d+1$ queries are necessary to achieve $\varepsilon(\delta) < 1$

  [any answers to any d queries agree with some codeword in $RS[\mathbb{F},L,d]$]

A query complexity of $O(d)$ is useful only when $d \ll |L|$.

But in our case $|L| = \Theta(d)$ and $d = \Theta(\text{computation size})$.

So we need query complexity $q \ll d$ (ideally $q = \text{poly}(\log d)$ or even $q = O(1)$).

What do we do?

The above considerations are for PROXIMITY TESTS.

We can ask the prover's help, which leads to a PROXIMITY PROOF.

Proximity Proofs

The prover helps the verifier to test proximity, e.g. to RS[𝔽,L,d].

One can consider DIFFERENT MODELS:

IPP (IP of proximity)

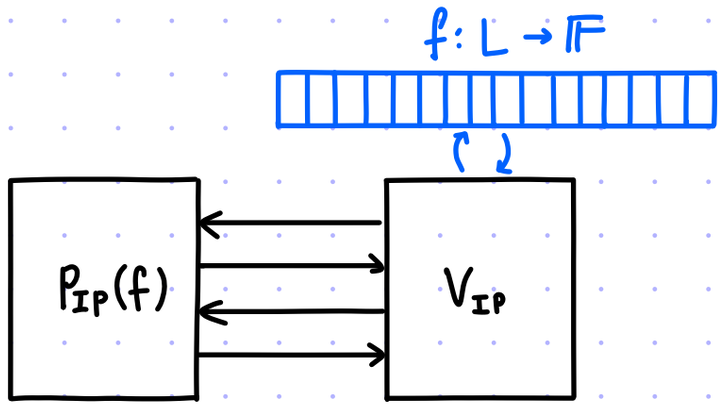

PCPP (PCP of proximity)

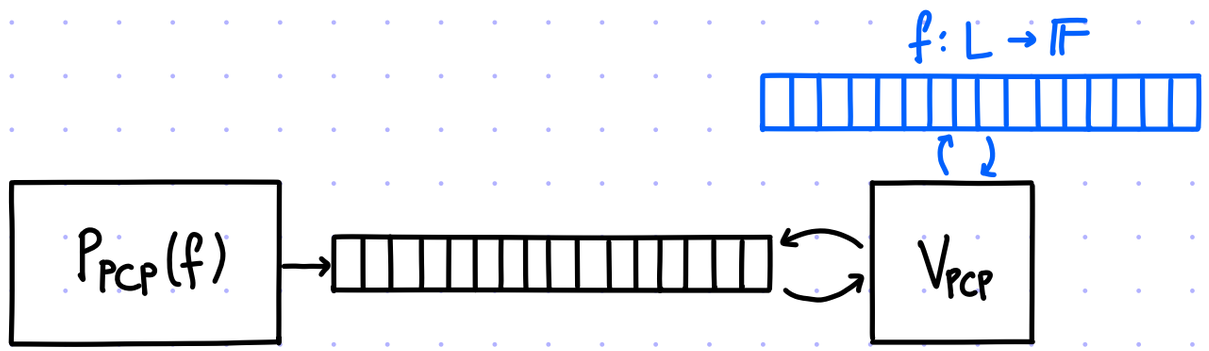

IOPP (IOP of proximity)

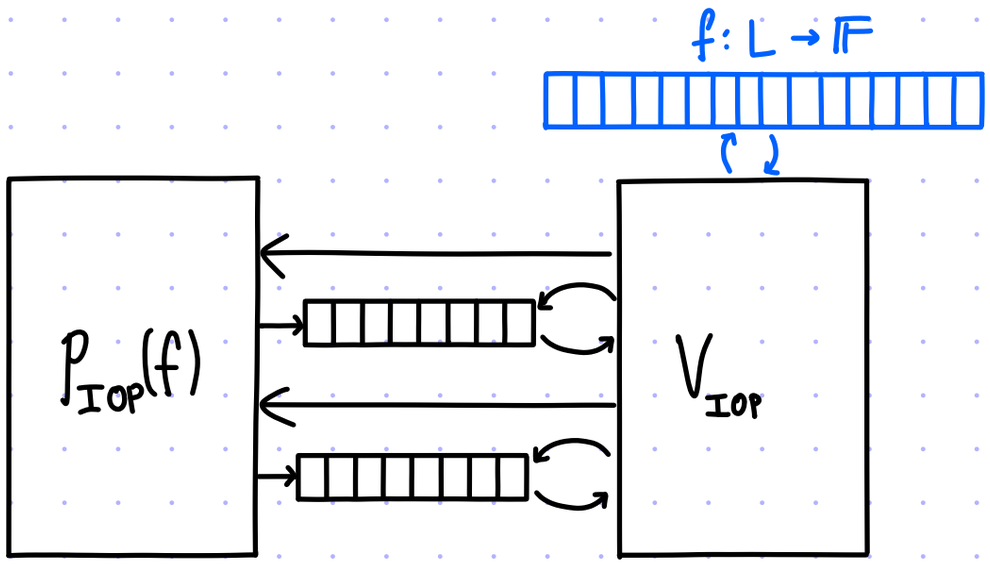

The verifier has oracle access to $f: L \to \mathbb{F}$.

The goal is to distinguish:

- $f \in RS[\mathbb{F},L,d]$
- f is far from $RS[\mathbb{F},L,d]$

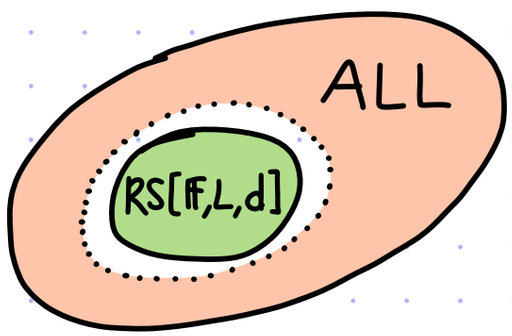

Today we study an IOPP for RS[𝔽,L,d].

Proximity Proofs for the Reed–Solomon Code

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

① COMPLETENESS: if $f \in RS[\mathbb{F},L,d]$ then $\Pr[\langle P(f), V^f \rangle = 1] = 1$.

② SOUNDNESS: if f is δ-far from $RS[\mathbb{F},L,d]$ then $\forall \tilde{P}$ $\Pr[\langle \tilde{P}, V^f \rangle = 1] \leq \varepsilon(\delta)$.

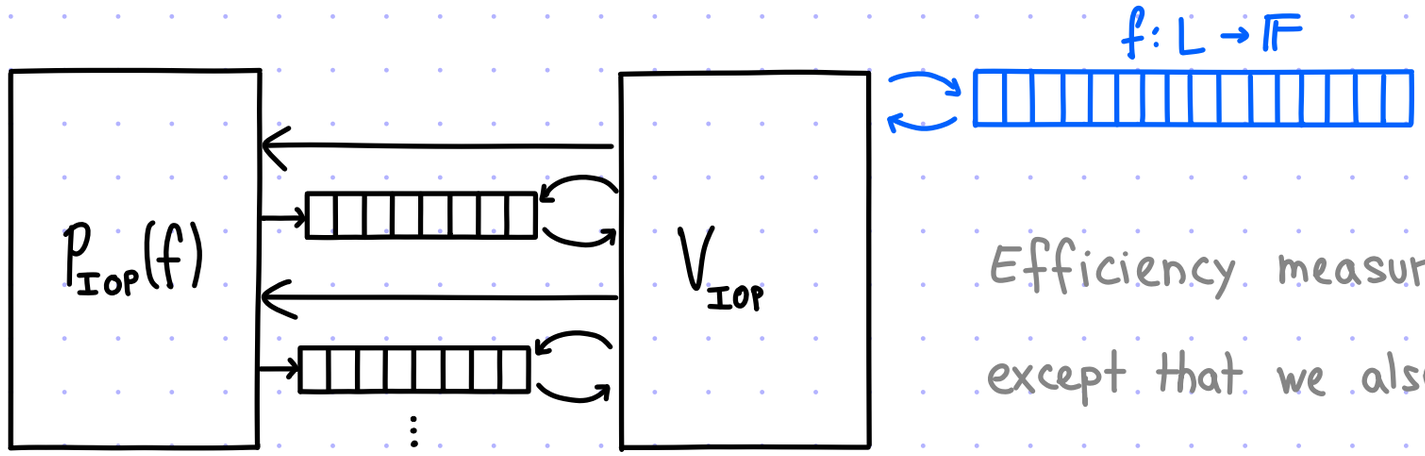

Efficiency measures are as in an IOP, except that we also charge for queries to f.

We restrict our attention to smooth domains: L is a subgroup of $\mathbb{F}^*$ with $|L| = 2^k$.

[Similar results hold for other (multiplicative or additive) subgroups L of $\mathbb{F}$.]

theorem: For every $\mathbb{F}$, smooth domain $L \subseteq \mathbb{F}$, and $d < |L|$,

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

This is called the FRI protocol. (FRI = Fast Reed-Solomon IOPP.)

This IOPP is important in practice. Its analysis raises elegant questions in coding theory.

Inspiration from the Fast Fourier Transform

We can write any polynomial $\hat{f} \in \mathbb{F}[x]$ as $\hat{g}(x^2) + x\hat{h}(x^2)$ where $\begin{cases} \hat{g} & \text{are the even coefficients} \\ \hat{h} & \text{are the odd coefficients} \end{cases}$.

The (radix-2) FFT uses a DIVIDE-AND-CONQUER approach:

Evaluate $\hat{f}(x)$ on $L = \langle \omega \rangle$:

1. Evaluate $\hat{g} := \text{even}(\hat{f})$ on $L^2 = \langle \omega^2 \rangle$.
2. Evaluate $\hat{h} := \text{odd}(\hat{f})$ on $L^2 = \langle \omega^2 \rangle$.
3. For $i = 0, 1, \ldots, \frac{|L|}{2} - 1$: $\hat{f}(\omega^i) := \hat{g}(\omega^{2i}) + \omega^i \cdot \hat{h}(\omega^{2i})$

(recall $-1 = \omega^{\frac{|L|}{2}}$ so $-\omega^i = \omega^{i+\frac{|L|}{2}}$) $\hat{f}(-\omega^i) := \hat{g}(\omega^{2i}) - \omega^i \cdot \hat{h}(\omega^{2i})$

(base case: if $\deg(\hat{f}) = 0$ then return $\hat{f}(1)$)

Diagram of $L = \langle \omega \rangle$ where $\omega = e^{\frac{2\pi i}{8}}$:

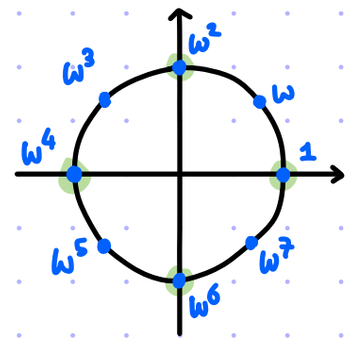

$L = \{1, \omega, \omega^2, \omega^3, \omega^4, \omega^5, \omega^6, \omega^7\}$

$L^2 = \{1, \omega^2, \omega^4, \omega^6\}$

$L^4 = \{1, \omega^4\}$

$L^8 = \{1\}$

The nested structure $L \supseteq L^2 \supseteq L^4 \supseteq \cdots$ enables recursion.

Both subproblems have half the size, and the recursion depth is $r = \log d$.

The total number of operations is $T(|L|) = 2 \cdot T(|L|/2) + O(|L|) = O(|L| \cdot \log |L|)$.

Back to low-degree testing:

- $f : L \to \mathbb{F}$ satisfies $\deg(\hat{f}) < d$
- $g : L^2 \to \mathbb{F}$ where $\hat{g} := \text{even}(\hat{f})$ satisfies $\deg(\hat{g}) < d/2$
- $h : L^2 \to \mathbb{F}$ where $\hat{h} := \text{odd}(\hat{f})$ satisfies $\deg(\hat{h}) < d/2$

Q: Can we devise a DIVIDE-AND-CONQUER approach to low-degree testing?

Attempt 1: Recurse on Each Subproblem

$P((\mathbb{F}, L, d), f)$

- Compute $\hat{g} := \text{even}(\hat{f})$, $\hat{h} := \text{odd}(\hat{f})$.
- Set $g := \hat{g}|_{L^2}$, $h := \hat{h}|_{L^2}$.

$f : L \to \mathbb{F}$

$g, h : L^2 \to \mathbb{F}$

recurse:

$g \stackrel{?}{\in} RS[\mathbb{F}, L^2, d/2]$

$h \stackrel{?}{\in} RS[\mathbb{F}, L^2, d/2]$

$V((\mathbb{F}, L, d))$

Sample $\mu \leftarrow L$ and check $f(\mu) \stackrel{?}{=} g(\mu^2) + \mu \cdot h(\mu^2)$.

PROBLEM: linear number of queries $(q(d) = 3 + 2 \cdot q(d/2) = \Theta(d))$

PROBLEM: distance decays in each recursion

Example:

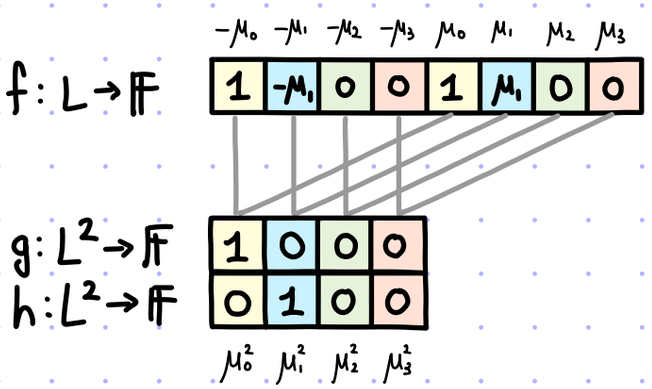

50% corrupted

25% corrupted

25% corrupted

The consistency check always passes:

$$\forall \mu \in L \quad f(\mu) = g(\mu^2) + \mu \cdot h(\mu^2)$$

Hence distance can drop as $\delta \to \delta/2 \to \delta/4 \to \cdots \to \delta/2^r$.

We cannot afford $r = \omega(1)$ rounds of interaction.

Attempt 2: Fold and Recurse [1/3]

$P((\mathbb{F},L,d),f)$

- Compute $\hat{g} := \text{even}(\hat{f})$, $\hat{h} := \text{odd}(\hat{f})$.
- Set $g := \hat{g}|_{L^2}$, $h := \hat{h}|_{L^2}$.
- Set $f_\alpha := g + \alpha \cdot h$.

$f : L \to \mathbb{F}$

$g, h : L^2 \to \mathbb{F}$ →

← α

$f_\alpha : L^2 \to \mathbb{F}$ →

recurse:

$f_\alpha \stackrel{?}{\in} RS[\mathbb{F}, L^2, d/2]$

$V((\mathbb{F},L,d))$

Sample $\mu \leftarrow L$ and check $f(\mu) \stackrel{?}{=} g(\mu^2) + \mu \cdot h(\mu^2)$.

Sample $\alpha \in \mathbb{F}$.

Check $f_\alpha(\mu^2) \stackrel{?}{=} g(\mu^2) + \alpha \cdot h(\mu^2)$.

The number of queries is $q(d) = 4 + q(d/2) = \Theta(\log d)$.

But does random folding work?

First we consider the NOISE-FREE CASE.

COMPLETENESS: if $\deg(\hat{f}) < d$ then $\deg(\hat{g}), \deg(\hat{h}) < d/2$ so $\forall \alpha \in \mathbb{F}$ $\deg(\hat{g} + \alpha \cdot \hat{h}) < d/2$

"SOUNDNESS": if $\deg(\hat{f}) \geq d$ then $\deg(\hat{g}) \geq d/2$ or $\deg(\hat{h}) \geq d/2$ so $\Pr_\alpha[\deg(\hat{g} + \alpha \cdot \hat{h}) \geq d/2] \geq 1 - \frac{1}{|\mathbb{F}|}$.

For $i := \max\{\deg(\hat{g}), \deg(\hat{h})\}$, $\Pr_\alpha[\deg(\hat{g} + \alpha \cdot \hat{h}) < \max\{\deg(\hat{g}), \deg(\hat{h})\}] = \Pr_\alpha[\text{coeff}(X^i, \hat{g}) + \alpha \cdot \text{coeff}(X^i, \hat{h}) = 0] \leq \frac{1}{|\mathbb{F}|}$.

Attempt 2: Fold and Recurse [2/3]

$P((\mathbb{F}, L, d), f)$

- Compute $\hat{g} := \text{even}(\hat{f})$, $\hat{h} := \text{odd}(\hat{f})$.
- Set $g := \hat{g}|_{L^2}$, $h := \hat{h}|_{L^2}$.
- Set $f_\alpha := g + \alpha \cdot h$.

$f : L \to \mathbb{F}$

$g, h : L^2 \to \mathbb{F}$ →

← α

$f_\alpha : L^2 \to \mathbb{F}$ →

recurse:

$f_\alpha \stackrel{?}{\in} RS[\mathbb{F}, L^2, d/2]$

$V((\mathbb{F}, L, d))$

Sample $\mu \leftarrow L$ and check $f(\mu) \stackrel{?}{=} g(\mu^2) + \mu \cdot h(\mu^2)$.

Sample $\alpha \in \mathbb{F}$.

Check $f_\alpha(\mu^2) \stackrel{?}{=} g(\mu^2) + \alpha \cdot h(\mu^2)$.

Next we consider the Noisy Case.

Example:

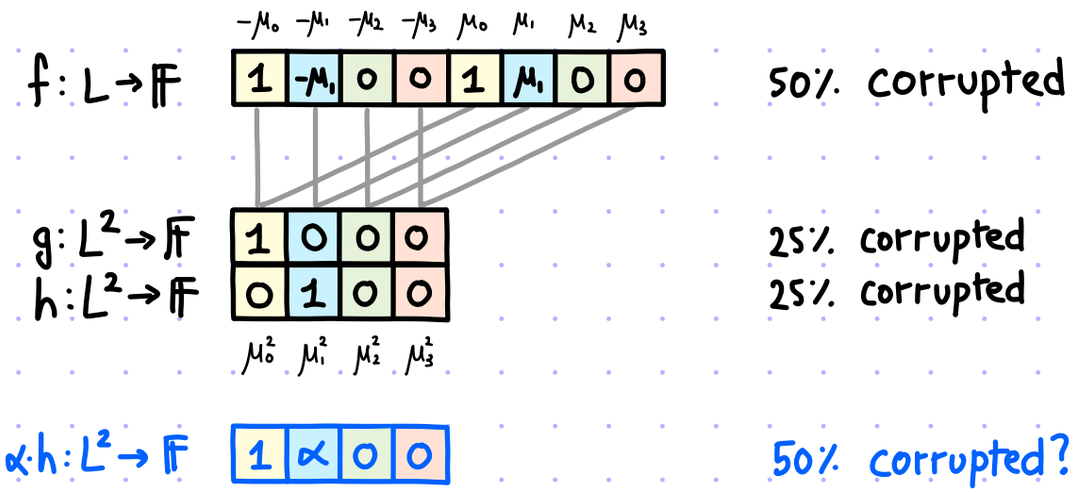

Random folding seems to preserve distance (at least whp).

Attempt 2: Fold and Recurse [3/3]

$P((\mathbb{F},L,d),f)$

- Compute $\hat{g} := \text{even}(\hat{f})$, $\hat{h} := \text{odd}(\hat{f})$.
- Set $g := \hat{g}|_{L^2}$, $h := \hat{h}|_{L^2}$.
- Set $f_\alpha := g + \alpha \cdot h$.

$f : L \to \mathbb{F}$

$g, h : L^2 \to \mathbb{F}$ →

← α

$f_\alpha : L^2 \to \mathbb{F}$ →

recurse: $f_\alpha \stackrel{?}{\in} RS[\mathbb{F}, L^2, d/2]$

$V((\mathbb{F},L,d))$

Sample $\mu \leftarrow L$ and check $f(\mu) \stackrel{?}{=} g(\mu^2) + \mu \cdot h(\mu^2)$.

Sample $\alpha \in \mathbb{F}$.

Check $f_\alpha(\mu^2) \stackrel{?}{=} g(\mu^2) + \alpha \cdot h(\mu^2)$.

What if a cheating prover sends inconsistent g, h, f_α?

There are consistency checks: between f and g, h; and between g, h and f_α.

Informally, in each round we pay:

ⓐ an error for the event that $g + \alpha \cdot h$ is close to RS even if f is far from RS;

ⓑ an error for the event that a consistency check fails.

There are $r = \Theta(\log d)$ rounds, so these errors should be $O(\frac{1}{r}) = O(\frac{1}{\log d})$.

The error in ⓐ is small. But the error in ⓑ is $\Theta(1)$.

What to do? [i] $O(\log\log d)$ queries/round (leads to $O(\log d \cdot \log\log d)$ total queries) [ii] improve the protocol.

The FRI Protocol

Changes from prior protocol:

- drop g and h (they are not needed)
- global multi-round consistency check

These changes lead to the FRI Protocol.

It is an IOPP for $RS[\mathbb{F},L,d] = \{ f: L \to \mathbb{F} : \deg(\hat{f}) < d \}$ where L is a subgroup of $\mathbb{F}^*$ with $|L| = 2^k$.

def: $\forall\, f: L \to \mathbb{F}$ $\forall\, \alpha \in \mathbb{F}$, define $Fold(f,\alpha): L^2 \to \mathbb{F}$ as $Fold(f,\alpha)(\gamma^2) := \frac{f(\gamma)+f(-\gamma)}{2} + \alpha \cdot \frac{f(\gamma)-f(-\gamma)}{2\cdot\gamma}$.

$P((\mathbb{F},L,d),f_0)$ $\quad$ $f_0: L \to \mathbb{F}$ $\quad$ $V((\mathbb{F},L,d))$

$\xleftarrow{\alpha_0 \in \mathbb{F}}$

$f_1 := Fold(f_0,\alpha_0)$ $\quad$ $\xrightarrow{f_1: L^2 \to \mathbb{F}}$

Interaction randomness: $\alpha_0, \alpha_1, \dots, \alpha_{r-1} \leftarrow \mathbb{F}$

Consistency check randomness: $\mu_1, \dots, \mu_t \leftarrow L$

$\xleftarrow{\alpha_1 \in \mathbb{F}}$

$f_2 := Fold(f_1,\alpha_1)$ $\quad$ $\xrightarrow{f_2: L^4 \to \mathbb{F}}$

For each repetition $j \in [t]$:

$\forall\, i \in \{0,1,\dots,r-1\}$

$f_{i+1}(\mu_j^{2^{i+1}}) \stackrel{?}{=} \frac{f_i(\mu_j^{2^i}) + f_i(-\mu_j^{2^i})}{2} + \alpha_i \cdot \frac{f_i(\mu_j^{2^i}) - f_i(-\mu_j^{2^i})}{2\mu_j^{2^i}}$

$\vdots$

$\xleftarrow{\alpha_{r-1} \in \mathbb{F}}$

$f_r := Fold(f_{r-1},\alpha_{r-1})$ $\quad$ $\xrightarrow{f_r: L^{2^r} \to \mathbb{F}}$

Low-degree check: $f_r \stackrel{?}{\in} RS[\mathbb{F}, L^{2^r}, d/2^r]$.

query pattern for $\mu \in L$

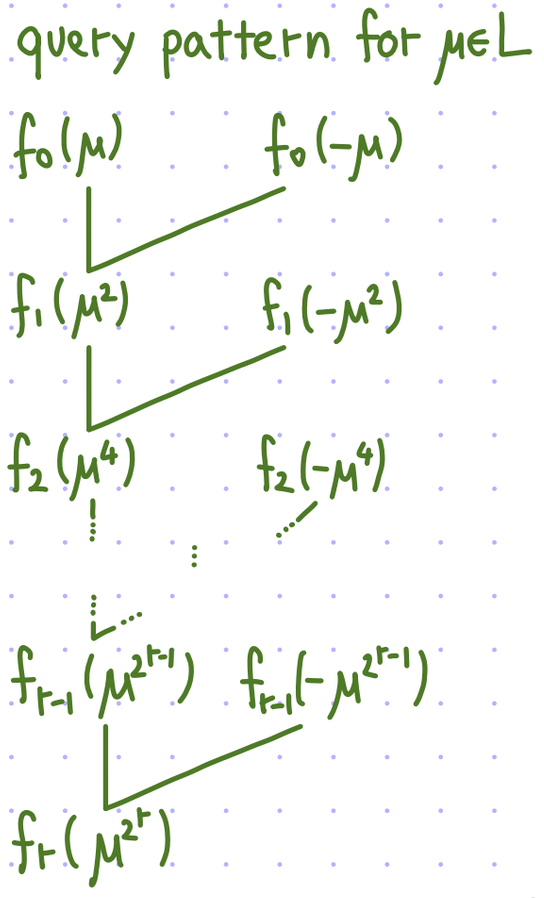

Completeness

lemma: FRI has perfect completeness

proof: Suppose that $f_0 \in RS[\mathbb{F}, L, d]$, so that $\deg(\hat{f}_0) < d$.

Fix interaction randomness:

$\alpha_0, \alpha_1, \ldots, \alpha_{r-1} \in \mathbb{F}$.

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

For every $i \in \{0,1,\ldots,r-1\}$, the i-th check passes $\forall\, \mu^{2^i} \in L^{2^i}$ by definition of $\text{Fold}(f_i, \alpha_i)$.

Moreover $f_0 \in RS[\mathbb{F}, L, d]$ implies $f_r \in RS[\mathbb{F}, L^{2^r}, d/2^r]$ so the degree check also passes. ■

claim: For every $i \in \{0,1,\ldots,r-1\}$: $f_i \in RS[\mathbb{F}, L^{2^i}, d/2^i] \to \forall \alpha_i \in \mathbb{F}$, $\text{Fold}(f_i, \alpha_i) \in RS[\mathbb{F}, L^{2^{i+1}}, d/2^{i+1}]$

proof: The hypothesis implies that $\deg(\hat{f}_i) < d/2^i$, so $\deg(\text{even}(\hat{f}_i)) < d/2^{i+1}$ and $\deg(\text{odd}(\hat{f}_i)) < d/2^{i+1}$.

Fix $\alpha_i \in \mathbb{F}$. By definition, $\deg(\widehat{\text{Fold}(f_i, \alpha_i)}) < |L^{2^{i+1}}|$. Also, $\deg(\text{even}(\hat{f}_i) + \alpha_i \,\text{odd}(\hat{f}_i)) < d/2^{i+1}$.

For every $\forall\, \gamma^2 \in L^{2^{i+1}}$, $\text{even}(\hat{f}_i)(\gamma^2) + \alpha_i \cdot \text{odd}(\hat{f}_i)(\gamma^2) = \frac{\hat{f}_i(\gamma) + \hat{f}_i(-\gamma)}{2} + \alpha_i \cdot \frac{\hat{f}_i(\gamma) - \hat{f}_i(-\gamma)}{2 \cdot \gamma} = \frac{f_i(\gamma) + f_i(-\gamma)}{2} + \alpha_i \cdot \frac{f_i(\gamma) - f_i(-\gamma)}{2 \cdot \gamma}$

$= \text{Fold}(f_i, \alpha_i)(\gamma^2) = \widehat{\text{Fold}(f_i, \alpha_i)}(\gamma^2)$. Hence $\widehat{\text{Fold}(f_i, \alpha_i)} \equiv \text{even}(\hat{f}_i) + \alpha_i \,\text{odd}(\hat{f}_i)$, so $\deg(\widehat{\text{Fold}(f_i, \alpha_i)}) < d/2^{i+1}$. ■

EFFICIENCY:

- prover time is $O(|L| + |L|/2 + |L|/4 + \cdots + |L|/2^{r-1}) = O(|L|)$
- verifier time is $O(t \cdot r + |L|/2^r) = O(t \cdot \log d)$ when $r = \log d$ and $|L| = \Theta(d)$
- query complexity is $O(t \cdot r + |L|/2^r) = O(t \cdot \log d)$ when $r = \log d$ and $|L| = \Theta(d)$

Intuition: A Simple Attack

$f_0: L \to \mathbb{F}$ ← $V((\mathbb{F}, L, d))$

$\alpha_0 \in \mathbb{F}$ → Interaction randomness: $\alpha_0, \alpha_1, \ldots, \alpha_{r-1} \leftarrow \mathbb{F}$

$f_1: L^2 \to \mathbb{F}$ → Consistency check randomness: $\mu_1, \ldots, \mu_t \leftarrow L$

$\alpha_1 \in \mathbb{F}$ → For each repetition $j \in [t]$:

$f_2: L^4 \to \mathbb{F}$ → $\forall\, i \in \{0, 1, \ldots, r-1\}$

$f_{i+1}(\mu_j^{2^{i+1}}) \stackrel{?}{=} \frac{f_i(\mu_j^{2^i}) + f_i(-\mu_j^{2^i})}{2} + \alpha_i \cdot \frac{f_i(\mu_j^{2^i}) - f_i(-\mu_j^{2^i})}{2\mu_j^{2^i}}$

$\alpha_{r-1} \in \mathbb{F}$ →

$f_r: L^{2^r} \to \mathbb{F}$ → Low-degree check: $f_r \stackrel{?}{\in} RS[\mathbb{F}, L^{2^r}, d/2^r]$.

We build intuition via a simple "attack".

claim: $\exists\, f_0: L \to \mathbb{F}$ that is δ-far from $RS[\mathbb{F}, L, d]$ and $\exists\, \tilde{P}$ s.t. $\Pr[\langle \tilde{P}, V^{f_0} \rangle = 1] \geq \max\left\{\frac{1}{|\mathbb{F}|}, (1-\delta)^t\right\}$.

proof: Split L into two sets L_0 and L_1 with

$|L_0| = (1-\delta) \cdot |L|$ and $|L_1| = \delta \cdot |L|$

Keeping elements with the same square together. (If $\mu \in L_b$ then $-\mu \in L_b$.)

Consider this $f_0: L \to \mathbb{F}$ that is δ-far from $RS[\mathbb{F}, L, d]$: $f_0(\mu) := \begin{cases} 0 & \text{if } \mu \in L_0 \\ g(\mu) & \text{if } \mu \in L_1 \end{cases}$

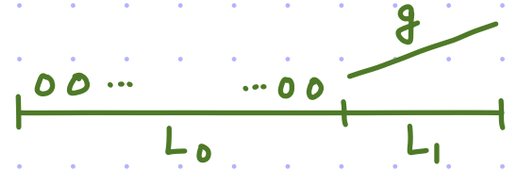

where $g(x) = ax + b$ is non-zero on L_1 and $a \neq 0$.

For every $\alpha_0 \in \mathbb{F}$, $Fold(f_0, \alpha_0)(\mu^2) = \begin{cases} 0 & \text{if } \mu \in L_0 \\ a\alpha_0 + b & \text{if } \mu \in L_1 \end{cases}$. Indeed: $\frac{g(\mu)+g(-\mu)}{2} + \alpha_0 \frac{g(\mu)-g(-\mu)}{2\mu} = b + \alpha_0 a$.

The prover $\tilde{P}$ sends $f_1, \ldots, f_r$ that (always) are zero functions.

- $\forall\, \alpha_0, \alpha_1, \ldots, \alpha_{r-1} \in \mathbb{F}$, $\Pr_{\mu_1, \ldots, \mu_t \in L}[\langle \tilde{P}, V^{f_0}(\vec{\alpha}, (\mu_1, \ldots, \mu_t)) \rangle = 1] \geq (1-\delta)^t$ ($\forall\, j \in [t]$, $\mu_j \in L_0 \to$ j-th consistency check passes)
- $\forall\, \mu_1, \ldots, \mu_t \in L$ $\forall\, \alpha_1, \ldots, \alpha_{r-1} \in \mathbb{F}$ $\Pr_{\alpha_0 \in \mathbb{F}}[\langle \tilde{P}, V^{f_0}(\vec{\alpha}, (\mu_1, \ldots, \mu_t)) \rangle = 1] = \frac{1}{|\mathbb{F}|}$ ($g(\alpha_0) = 0 \to \forall\, \mu \in L$ $Fold(f_0, \alpha_0)(\mu^2) = 0$)

NOTE: the same idea works with $f_0(\mu) := \begin{cases} p(\mu) & \text{if } \mu \in L_0 \\ g(\mu) & \text{if } \mu \in L_1 \end{cases}$ where $\deg(p) < d$ and $\forall \mu \in L_1$ $p(\mu) \neq g(\mu)$. Hence the case $p \equiv 0$ is not special.

Soundness

We saw a lower bound on soundness error:

claim: $\exists\ f_0 : L \to \mathbb{F}$ that is δ-far from $RS[\mathbb{F}, L, d]$ and $\exists\ \tilde{P}$ s.t. $\Pr[\langle \tilde{P}, V^{f_0} \rangle = 1] \geq \max\left\{\frac{1}{|\mathbb{F}|}, (1-\delta)^t\right\}$.

Here is an upper bound:

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

$$\forall \tilde{P},\ \Pr[\langle \tilde{P}, V^{f(t)} \rangle = 1] \leq O\left(\frac{|L|}{|\mathbb{F}|}\right) + \left(1 - \min\left\{\delta, c\left(\frac{d}{|L|}\right)\right\}\right)^t$$

$c\left(\frac{d}{|L|}\right)$ is a universal constant that depends on the rate $d/|L|$

$f_0 : L \to \mathbb{F}$

$\xleftarrow{\alpha_0 \in \mathbb{F}}$

$\xrightarrow{f_1 : L^2 \to \mathbb{F}}$

$\xleftarrow{\alpha_1 \in \mathbb{F}}$

$\xrightarrow{f_2 : L^4 \to \mathbb{F}}$

$\vdots$

$\xleftarrow{\alpha_{r-1} \in \mathbb{F}}$

$\xrightarrow{f_r : L^{2^r} \to \mathbb{F}}$

$V((\mathbb{F}, L, d))$

Interaction randomness: $\alpha_0, \alpha_1, \ldots, \alpha_{r-1} \leftarrow \mathbb{F}$

Consistency check randomness: $\mu_1, \ldots, \mu_t \leftarrow L$

For each repetition $j \in [t]$:

$\forall\ i \in \{0, 1, \ldots, r-1\}$

$$f_{i+1}(\mu_j^{2^{i+1}}) \stackrel{?}{=} \frac{f_i(\mu_j^{2^i}) + f_i(-\mu_j^{2^i})}{2} + \alpha_i \cdot \frac{f_i(\mu_j^{2^i}) - f_i(-\mu_j^{2^i})}{2\mu_j^{2^i}}$$

Low-degree check: $f_r \stackrel{?}{\in} RS[\mathbb{F}, L^{2^r}, d/2^r]$.

More precisely, $\Pr_{\alpha_0, \ldots, \alpha_{r-1}}\left[\Pr_{\mu_1, \ldots, \mu_t}\left[\langle \tilde{P}, V^{f(t;(\vec{\alpha}, \vec{\mu}))} \rangle = 1\right] \leq \left(1 - \min\left\{\delta, c\left(\frac{d}{|L|}\right)\right\}\right)^t\right] \geq 1 - \Omega\left(\frac{|L|}{|\mathbb{F}|}\right)$.

We prove the theorem in the next lecture.

The proof relies on fundamental statements about WORST-VS-AVERAGE-CASE DISTANCES TO SUBSPACES.

KNOWN: tighter upper bounds (using algebraic geometry and algebraic function fields)

OPEN: tight soundness analysis (it would lead to efficiency gains in practice)

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