Lecture 17

Foundations of Probabilistic Proofs
Fall 2020
Alessandro Chiesa

Proximity Testing to the Reed--Solomon Code

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

① completeness: if $f \in RS[\mathbb{F}, L, d]$ then the test accepts w.p. 1

② soundness: if f is δ-far from $RS[\mathbb{F}, L, d]$ then the test accepts w.p. $\leq \varepsilon(\delta)$ (with $\delta = \Omega(1)$ $\to \varepsilon(\delta) = O(1)$)

We have seen that:

- $d+2$ queries suffice to achieve $\varepsilon(\delta) = 1 - \delta$

  [interpolate the answers to any $d+1$ queries, and check consistency with the answer to a random query]

- $d+2$ queries are necessary to achieve $\varepsilon(\delta) < 1$

  [any answers to any $d+1$ queries are consistent with some codeword in $RS[\mathbb{F}, L, d]$]

This is ok when $d \ll |L|$ but in our case $d = \Theta(n)$ and $|L| = \Theta(d)$,

so we need query complexity that is much less than d (ideally, $\mathrm{poly}(\log d)$ or $O(1)$).

What do we do?

The above considerations are about proximity tests only.

We have the option of asking the prover's help, which leads us to a proximity proof.

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 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 $\forall \tilde{P}$ $\Pr[\langle \tilde{P}, V^f\rangle = 1] \leq \varepsilon(\delta)$

An IOPP for RS look like this:

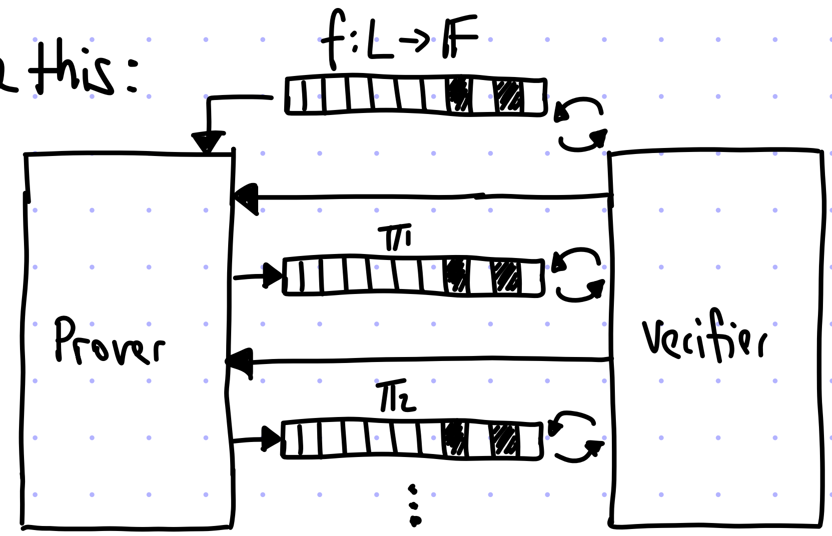

The efficiency measures are as in an IOP except we also charge for queries to f.

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

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

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

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.]

Inspiration from the Fast Fourier Transform

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

The (radix-2) FFT is based on the following divide-and-conquer approach:

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

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

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

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

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

Back to low-degree testing: $f : L \to \mathbb{F}$, $\deg(\hat{f}) \circledast< d$ iff $g, h : L^2 \to \mathbb{F}$, $\deg(\hat{g}), \deg(\hat{h}) \circledast< d/2$ $\left[\begin{array}{l} \text{for } \hat{g} := \mathrm{even}(\hat{f}) \\ \& \ \hat{h} := \mathrm{odd}(\hat{f}) \end{array}\right]$

Can we devise a divide-and-conquer approach to low-degree testing?

$\circledast$ for the rest of today we use strictly less than d

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 $f(\mu) \stackrel{?}{=} g(\mu^2) + \mu h(\mu^2)$

recurse to test that
$g \in RS[\mathbb{F}, L^2, d/2]$
$h \in RS[\mathbb{F}, L^2, d/2]$

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

Problem: it's not even a test because distance decays in each recursion

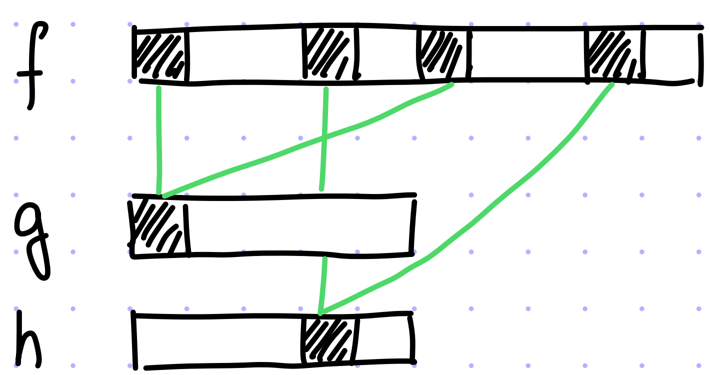

40 errors out of 100 locations $\delta = 40\%$

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

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

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

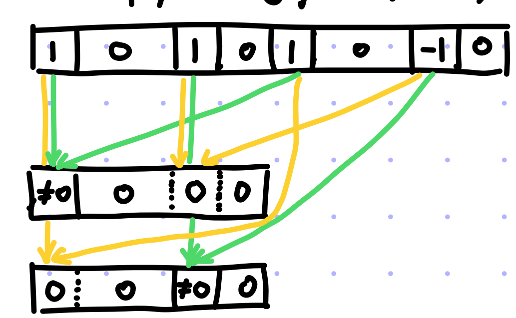

The distance could drop as $\delta \to \delta/2 \to \delta/4 \to \ldots \to \delta/2^r$.

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

Attempt 2: Fold and Recurse

$P((\mathbb{F},L,d),f)$	$f: L\to\mathbb{F}$	$V((\mathbb{F},L,d))$
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 $f(\mu) \stackrel{?}{=} g(\mu^2)+\mu h(\mu^2)$
	$\xleftarrow{\alpha}$	sample $\alpha\in\mathbb{F}$
Set $f_\alpha := g+\alpha h$	$\xrightarrow{f_\alpha:\, L^2\to\mathbb{F}}$	check $f_\alpha(\mu^2) = g(\mu^2)+\alpha h(\mu^2)$
	recourse to test that $f_\alpha \in RS[\mathbb{F},L^2,d/2]$	

The number of queries is now $q(d) = 4 + q(d/2) = O(\log d)$. This is good.

But does random folding make sense?

Let's consider the noise-free case first:

- completeness: if $\deg(\hat{f})<d$ then $\deg(\hat{h}),\deg(\hat{g})<d/2$ so $\forall\alpha\in\mathbb{F}$ $\deg(\hat{g}+\alpha\hat{h})<d/2$
- "soundness": if $\deg(\hat{f})\geq d$ then either $\deg(\hat{h})\geq d/2$ or $\deg(\hat{g})\geq d/2$,
  in which case $\Pr_\alpha\left[\deg(\hat{g}+\alpha\hat{h})\geq d/2\right]\geq 1-\frac{1}{|\mathbb{F}|}$

Indeed $\Pr_\alpha\left[\deg(\hat{g}+\alpha\hat{h})<\max\{\deg(\hat{g}),\deg(\hat{f})\}\right]\leq\frac{1}{|\mathbb{F}|}$ as there is 1 choice of α for which the highest-degree monomial is not in $\hat{g}+\alpha\hat{f}$

Attempt 2: Fold and Recurse

$P((\mathbb{F},L,d),f)$	$f: L\to\mathbb{F}$	$V((\mathbb{F},L,d))$
Compute $\hat{g} := \text{even}(\hat{f})$ and $\hat{h} := \text{odd}(\hat{f})$ and set $g := \hat{g}\rvert_{L^2}$ and $h := \hat{h}\rvert_{L^2}$	$\xrightarrow{g,h:\, L^2\to\mathbb{F}}$	sample $\mu\in L$ and check $f(\mu) \stackrel{?}{=} g(\mu^2)+\mu h(\mu^2)$
	$\xleftarrow{\alpha}$	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 $RS[\mathbb{F},L,d]$.

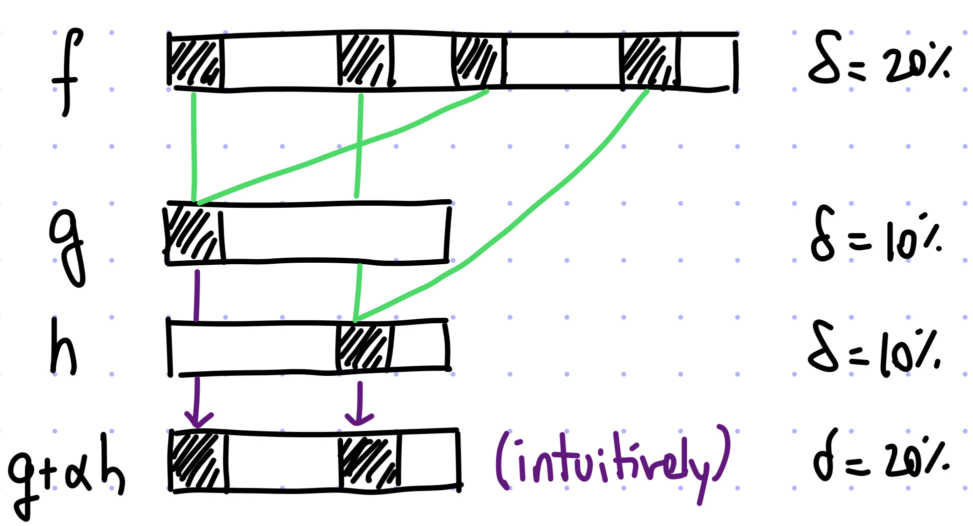

Folding seems to address the prior problem by preserving distance!

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

Two changes from prior protocol: drop g, h as they are not needed; do a single multi-round consistency check.

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

<u>lemma</u>: $\widehat{\mathrm{Fold}(f,\alpha)}(X) \equiv \mathrm{even}(\hat{f})(X) + \alpha\, \mathrm{odd}(\hat{f})(X)$

<u>proof</u>: For every $\gamma^2 \in L^2$, $\mathrm{even}(\hat{f})(\gamma^2) + \alpha \cdot \mathrm{odd}(\hat{f})(\gamma^2) = \frac{\hat{f}(\gamma)+\hat{f}(-\gamma)}{2} + \alpha \frac{\hat{f}(\gamma)-\hat{f}(-\gamma)}{2\gamma} = \widehat{\mathrm{Fold}(f,\alpha)}(\gamma^2)$.

These changes lead to the FRI protocol:

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

$P((\mathbb{F},L,d),f_0)$		$V((\mathbb{F},L,d))$
	$\xleftarrow{\alpha_0}$	
$f_1 := \mathrm{Fold}(f_0,\alpha_0)$	$\xrightarrow{f_1: L^2 \to \mathbb{F}}$	
	$\xleftarrow{\alpha_1}$	
$f_2 := \mathrm{Fold}(f_1,\alpha_1)$	$\xrightarrow{f_2: L^4 \to \mathbb{F}}$	
	$\vdots$	
	$\xleftarrow{\alpha_{r-1}}$	
$f_r := \mathrm{Fold}(f_{r-1},\alpha_{r-1})$	$\xrightarrow{f_r: L^{2^r} \to \mathbb{F}}$	

Verifier:

- interaction randomness: $\alpha_0, \alpha_1, \ldots, \alpha_{r-1} \leftarrow \mathbb{F}$
- consistency check randomness: $\mu \leftarrow L$ [can repeat t times]
- consistency check on $f_0, f_1, \ldots, f_r$:

$$f_1(\mu^2) \stackrel{?}{=} \frac{f_0(\mu)+f_0(-\mu)}{2} + \alpha_0 \cdot \frac{f_0(\mu)-f_0(-\mu)}{2\mu}$$

$$f_2(\mu^4) \stackrel{?}{=} \frac{f_1(\mu^2)+f_1(-\mu^2)}{2} + \alpha_1 \cdot \frac{f_1(\mu^2)-f_1(-\mu^2)}{2\mu^2}$$

$$\vdots$$

$$f_r(\mu^{2^r}) \stackrel{?}{=} \frac{f_{r-1}(\mu^{2^{r-1}})+f_{r-1}(-\mu^{2^{r-1}})}{2} + \alpha_{r-1} \cdot \frac{f_{r-1}(\mu^{2^{r-1}})-f_{r-1}(-\mu^{2^{r-1}})}{2\mu^{2^{r-1}}}$$

- last function is low degree: $f_r \stackrel{?}{\in} RS[\mathbb{F}, L^{2^r}, d/2^r]$

query pattern:

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

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

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

$\vdots$

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

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

Completeness

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

claim: FRI has perfect completeness

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

Fix any choice of interaction randomness:

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

For every $i = 1,\dots,r$, define $\hat{f}_i(X) := \mathrm{even}(\hat{f}_{i-1})(X) + \alpha_{i-1} \cdot \mathrm{odd}(\hat{f}_{i-1})(X)$.

Since $\deg(\hat{f}_0) < d$ we know that $\deg(\hat{f}_i) < d/2^i$ and thus $f_i := \hat{f}_i|_{L^{2^i}} \in RS[\mathbb{F}, L^{2^i}, d/2^i]$.

Observe that $f_i = \mathrm{Fold}(f_{i-1},\alpha_{i-1})$ because

$$\forall \gamma \in L^{2^{i-1}} \quad f_i(\gamma^2) = \mathrm{even}(\hat{f}_{i-1})(\gamma^2) + \alpha_{i-1} \cdot \mathrm{odd}(\hat{f}_{i-1})(\gamma^2) = \frac{f_{i-1}(\gamma)+f_{i-1}(-\gamma)}{2} + \alpha_{i-1} \cdot \frac{f_{i-1}(\gamma)-f_{i-1}(-\gamma)}{2\gamma}.$$

Hence for every $\mu \in L$ all the verifier consistency checks pass.

Finally, $f_r \in RS[\mathbb{F}, L^{2^r}, d/2^r]$ as argued above, so the verifier's degree check also passes. ∎

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

Intuition: A Simple Attack

- $f_0 : L \to \mathbb{F}$
- $\xleftarrow{\alpha_0}$
- $f_1 : L^2 \to \mathbb{F}$ $\longrightarrow$
- $\xleftarrow{\alpha_1}$
- $f_2 : L^4 \to \mathbb{F}$ $\longrightarrow$
- $\vdots$
- $\xleftarrow{\alpha_{r-1}}$
- $f_r : L^{2^r} \to \mathbb{F}$ $\longrightarrow$

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

- interaction randomness: $\alpha_0, \alpha_1, \dots, \alpha_{r-1} \leftarrow \mathbb{F}$
- consistency check randomness: $\mu \leftarrow L$ [can repeat t times]
- consistency check on $f_0, f_1, \dots, f_r$:

$$f_1(\mu^2) \stackrel{?}{=} \frac{f_0(\mu)+f_0(-\mu)}{2} + \alpha_0 \cdot \frac{f_0(\mu)-f_0(-\mu)}{2\mu}$$

$$f_2(\mu^4) \stackrel{?}{=} \frac{f_1(\mu^2)+f_1(-\mu^2)}{2} + \alpha_1 \cdot \frac{f_1(\mu^2)-f_1(-\mu^2)}{2\mu^2}$$

$$\vdots$$

$$f_r(\mu^{2^r}) \stackrel{?}{=} \frac{f_{r-1}(\mu^{2^{r-1}})+f_{r-1}(-\mu^{2^{r-1}})}{2} + \alpha_{r-1} \cdot \frac{f_{r-1}(\mu^{2^{r-1}})-f_{r-1}(-\mu^{2^{r-1}})}{2\mu^{2^{r-1}}}$$

- last function is low degree: $f_r \stackrel{?}{\in} RS[\mathbb{F}, L^{2^r}, d/2^r]$

Let's build intuition via a simple "attack".

claim: there is a prover strategy to make the verifier accept some δ-far f_0 w.p. $\geq \max\left\{\frac{1}{|\mathbb{F}|}, (1-\delta)^t\right\}$

dependence on field size necessary (and is a low-order term for large fields)

can view as the probability that all t queries to f_0 don't see noise

proof: Split L into two sets L_0 and L_1 with $|L_0| = (1-\delta)\cdot|L|$ and $|L_1| = \delta \cdot |L|$, while also keeping elements with the same square in the same set. (If $\mu \in L_b$ then $-\mu \in L_b$.)

Consider this $f_0 : L \to \mathbb{F}$ that is δ-far from $RS[\mathbb{F}, L, d]$:

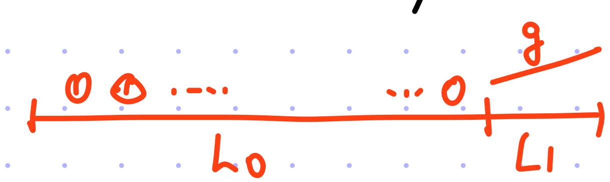

Here g is any line with no zeros on L_1.

Fix all other $f_2, f_3, \dots, f_r$ to be the zero functions.

Observe that:

- $\forall \alpha_0, \alpha_1, \dots, \alpha_{r-1} \in \mathbb{F}$ $\Pr_{\vec{\mu}}[\langle P, V^{f_0}(\alpha, \vec{\mu})\rangle = 1] \geq (1-\delta)^t$ (if $\vec{\mu} \in L_0^t$ then consistency tests pass)
- $\forall \vec{\mu} \in L^t$ $\forall \alpha_1, \dots, \alpha_{r-1} \in \mathbb{F}$ $\Pr_{\alpha_1}[\langle P, V^{f_0}(\alpha, \vec{\mu})\rangle = 1] = \frac{1}{|\mathbb{F}|}$ (if α_0 is a root of line)

[Note that 0 is not special: we can shift all the zeros to any low-degree polynomial.]

Soundness

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

$\xleftarrow{\ \alpha_0\ }$

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

$\xleftarrow{\ \alpha_1\ }$

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

$\vdots$

$\xleftarrow{\ \alpha_{r-1}\ }$

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

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

- interaction randomness: $\alpha_0, \alpha_1, \dots, \alpha_{r-1} \leftarrow \mathbb{F}$
- consistency check randomness: $\mu \leftarrow L$ [can repeat t times]
- consistency check on $f_0, f_1, \dots, f_r$:

$$f_1(\mu^2) \stackrel{?}{=} \frac{f_0(\mu)+f_0(-\mu)}{2} + \alpha_0 \cdot \frac{f_0(\mu)-f_0(-\mu)}{2\mu}$$

$$f_2(\mu^4) \stackrel{?}{=} \frac{f_1(\mu^2)+f_1(-\mu^2)}{2} + \alpha_1 \cdot \frac{f_1(\mu^2)-f_1(-\mu^2)}{2\mu^2}$$

$$\vdots$$

$$f_r(\mu^{2^r}) \stackrel{?}{=} \frac{f_{r-1}(\mu^{2^{r-1}})+f_{r-1}(-\mu^{2^{r-1}})}{2} + \alpha_{r-1} \cdot \frac{f_{r-1}(\mu^{2^{r-1}})-f_{r-1}(-\mu^{2^{r-1}})}{2\mu^{2^{r-1}}}$$

- last function is low degree: $f_r \stackrel{?}{\in} RS[\mathbb{F}, L^{2^r}, d/2^r]$

We have seen this lower bound on soundness error:

claim: there is a prover strategy to make the verifier accept some δ-far f_0 w.p. $\geq \max\left\{\frac{1}{|\mathbb{F}|}, (1-\delta)^t\right\}$

The upper bound is, to a first order, very close:

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

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

Here $c\left(\frac{d}{|L|}\right)$ is a universal constant with a dependence on the rate $d/|L|$.

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

We prove the theorem in the next lecture.

The proof relies on fundamental statements about worst-case vs average-case distances to subspaces.

Tighter upper bounds are known (which rely on tools from algebraic geometry and algebraic function fields), which lead to more efficiency in practice.

A tight soundness analysis remains an exciting open problem!