FRI and Proximity Proofs: continued

Dan Boneh
Stanford University

Review: Poly-IOP ⇒ IOP ⇒ SNARK

A direct SNARK construction:

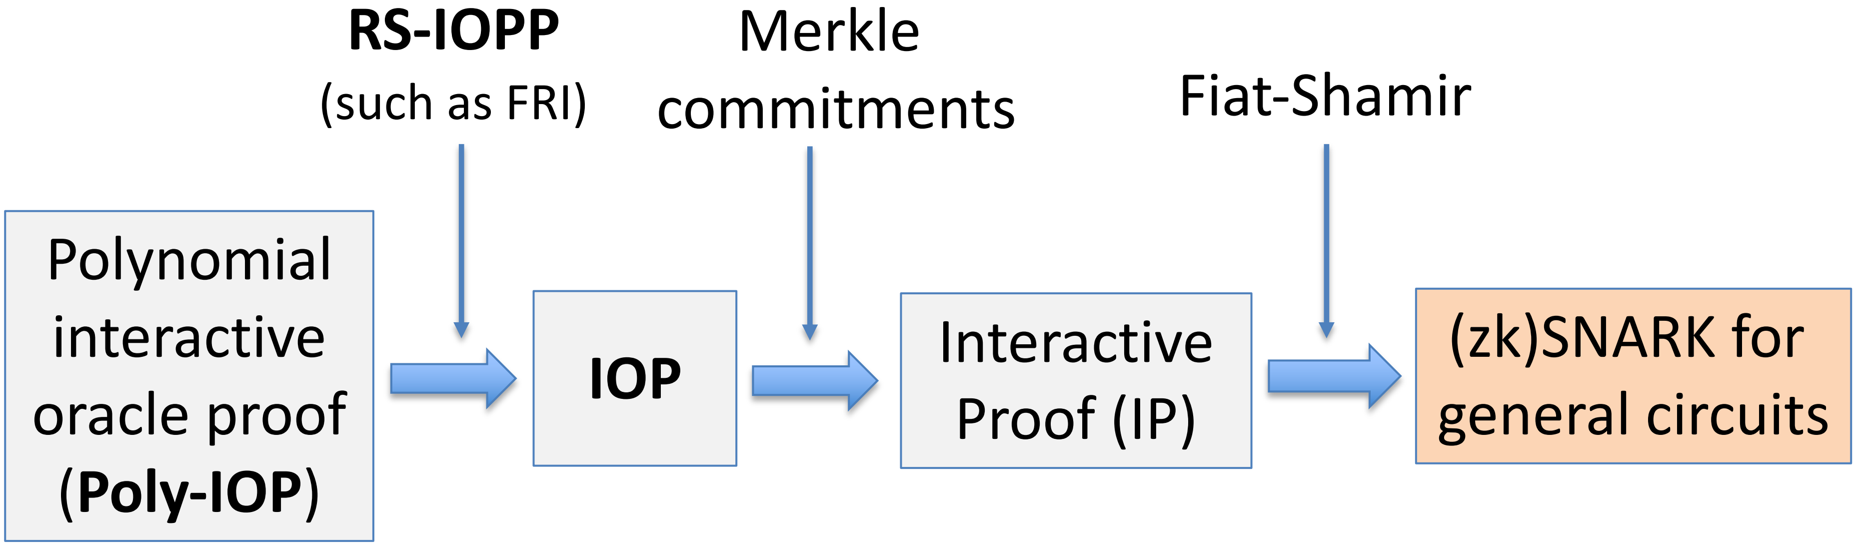

An RS-IOPP is the key ingredient in compilation

FRI: a Reed-Solomon IOP of Proximity

Let $\mathcal{C} = \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$, $u_0 \colon \mathcal{L} \to \mathbb{F}$, and $\delta \in [0,1]$

Prover $\boldsymbol{P}(\mathcal{C}, u_0, \cdot)$

Verfier $\boldsymbol{V}^{u_0}(\mathcal{C})$

Goal:

- $u_0 \in \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ $\Rightarrow$ Verifier outputs accept (with prob. 1)
- u_0 is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ $\Rightarrow$ Verifier outputs reject w.h.p

We don't care what happens when u_0 is between the two cases

Why needed? A key tool for compiling a Poly-IOP into an IOP.

FRI: a Reed-Solomon IOP of Proximity

Let $\mathcal{C} = \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$, $u_0 : \mathcal{L} \to \mathbb{F}$, and $\delta \in [0,1]$

<u>Prover $\boldsymbol{P}(\mathcal{C}, u_0, \cdot)$</u>

<u>Verfier $\boldsymbol{V}^{u_0}(\mathcal{C})$</u>

We will set:

$\mathcal{L} = \{1, \omega, \omega^2, \dots, \omega^{n-1}\} \subseteq \mathbb{F}$ and n, d are powers of two,

where $\omega^n = 1$ is an n-th primitive root of unity

Then: $|\mathcal{L}^2| \coloneqq |\{a^2 : a \in \mathcal{L}\}| = |\mathcal{L}|/2 = n/2$ $\quad (-a,\ a \to a^2)$

$|\mathcal{L}^4| \coloneqq |\{a^4 : a \in \mathcal{L}\}| = |\mathcal{L}|/4 = n/4$

Distance Preserving Transformations

Towards an efficient RS-IOPP

Distance Preserving Transformations

Let $\mathcal{L}, \mathcal{L}' \subseteq \mathbb{F}$, d, d' some degree bounds, and $\delta \in [0,1]$.

Def: A **distance preserving transformation** is a randomized map

$$T(u_1, \ldots, u_k; r) \to u$$

that maps $u_1, \ldots, u_k : \mathcal{L} \to \mathbb{F}$ to $u : \mathcal{L}' \to \mathbb{F}$ such that:

case 1: (the honest case)

if $u_1, \ldots, u_k \in \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ then $u \in \mathrm{RS}[\mathbb{F}, \mathcal{L}', d']$ for all r.

case 2: (the dishonest case)

if some u_j is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ then
u is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}', d']$, w.h.p over r.

Example 1: batch RS-IOPP

Setting: Prover has $u_0, \ldots, u_k : \mathcal{L} \to \mathbb{F}$, Verifier has oracles for $u_0, \ldots, u_k$.

Goal: convince Verifier that all $u_0, \ldots, u_k$ are δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$.

- **Naively**: run k RS-IOPP protocols $\Rightarrow$ expensive
- **Better**: batch all k into a single function $u: \mathcal{L} \to \mathbb{F}$

step 1: Verifier samples random r in $\mathbb{F}$; sends to prover

step 2: Prover sets $u \coloneqq u_0 + r \cdot u_1 + r^2 u_2 + \cdots + r^k u_k : \mathcal{L} \to \mathbb{F}$

step 3: Both run RS-IOPP on $u: \mathcal{L} \to \mathbb{F}$

when Verifier wants $u(a)$ for some $a \in \mathcal{L}$, prover opens all $u_0(a), \ldots, u_k(a)$

Why is this distance preserving?

Case 1: (an honest prover)

if $u_0, \ldots, u_k \in \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ then $u \in \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ for all $r \in \mathbb{F}$

Case 2: (a dishonest prover)

if some u_j is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$, we need to argue that u is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$, with high probability over $r \in \mathbb{F}$

When $\delta \in [0, 1 - \sqrt{\rho})$, Case 2 follows from the celebrated BCIKS proximity gap theorem.

The proximity gap theorem

Thm (BCIKS'20, Thm. 6.2): $\mathrm{RS}[\mathbb{F},\mathcal{L},d]$ an RS-code with const. rate $\rho \coloneqq d/n$ (say, $\rho = 0.5$)

Let $u_0, \ldots, u_k : \mathcal{L} \to \mathbb{F}$ and $0 < \delta < 1 - 1.01\sqrt{\rho}$. $n \coloneqq |\mathcal{L}|$

For $r \in \mathbb{F}$ define $u^{(r)} \coloneqq u_0 + r \cdot u_1 + r^2 u_2 + \cdots + r^k u_k$.

Suppose that $\Pr_r[\ u^{(r)}$ is δ-close to $\mathrm{RS}[\mathbb{F},\mathcal{L},d]\ ] > err$

then all u_j are δ-close to $\mathrm{RS}[\mathbb{F},\mathcal{L},d]$,

where
$$\begin{cases} err = O\left(\frac{kn}{|\mathbb{F}|}\right) & \text{for } 0 < \delta < \frac{1-\rho}{2} \\ err = O\left(\frac{kn^2}{|\mathbb{F}|}\right) & \text{for } \frac{1-\rho}{2} < \delta < 1 - 1.01\sqrt{\rho} \end{cases}$$

We will assume that err is negligible, i.e. err $< 1/2^{128}$
(if not, use multiple r)

The proximity gap theorem

Suppose that $\Pr_r[\ u^{(r)}$ is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]\ ] > err$
then all u_j are δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$

Contra-positive: if some u_j is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$
then $u^{(r)}$ is δ-far with high probability, over r.

Proximity gap error (err) as a function of $\delta \in [0,1]$:

conjectured to be small

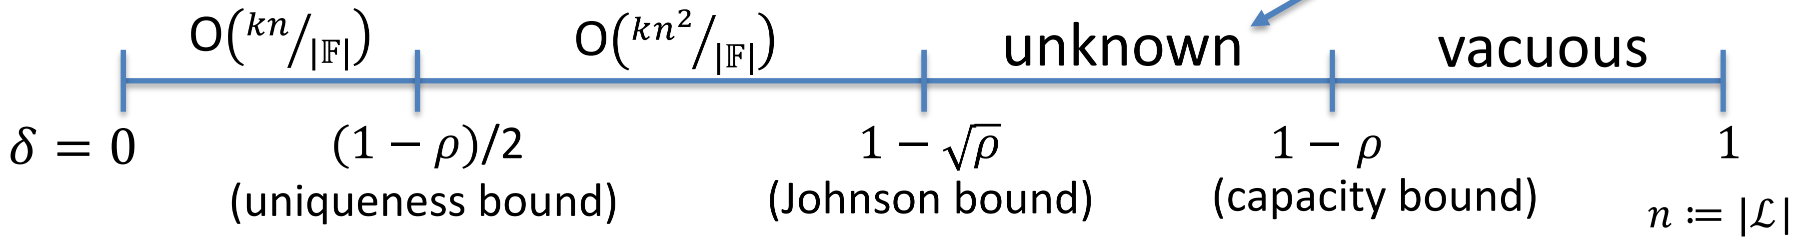

A stronger form: correlated proximity

Thm (BCIKS'20, Thm. 6.2):

Let $u_0, \ldots, u_k : \mathcal{L} \to \mathbb{F}$ and $0 < \delta < 1 - 1.01\sqrt{\rho}$.

Suppose that $\Pr_r[\ u^{(r)}$ is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]\ ] > err$

then there is an $S \subseteq \mathcal{L}$ such that $|S| \geq (1 - \delta) \cdot |\mathcal{L}|$ and

for all j: $\exists f_j \in \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ s.t. $\forall x \in S$: $u_j(x) = f_j(x)$

$\Rightarrow$ $u_0, \ldots, u_k$ are δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ on the same positions S .

(recall $u^{(r)} := u_0 + r \cdot u_1 + r^2 u_2 + \cdots + r^k u_k$)

Why is this called a proximity gap??

Suppose that $\Pr_r[\ u^{(r)}$ is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]\ ] > err$ then all u_j are δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ on the same positions $S \subseteq \mathcal{L}$

But if all $u_0, \ldots, u_k : \mathcal{L} \to \mathbb{F}$ are δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ on positions $S \subseteq \mathcal{L}$, then $u^{(r)}$ is δ-close for all $r \in \mathbb{F}$.

So $\Pr_r[\ u^{(r)}$ is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]\ ]$ exhibits a gap:

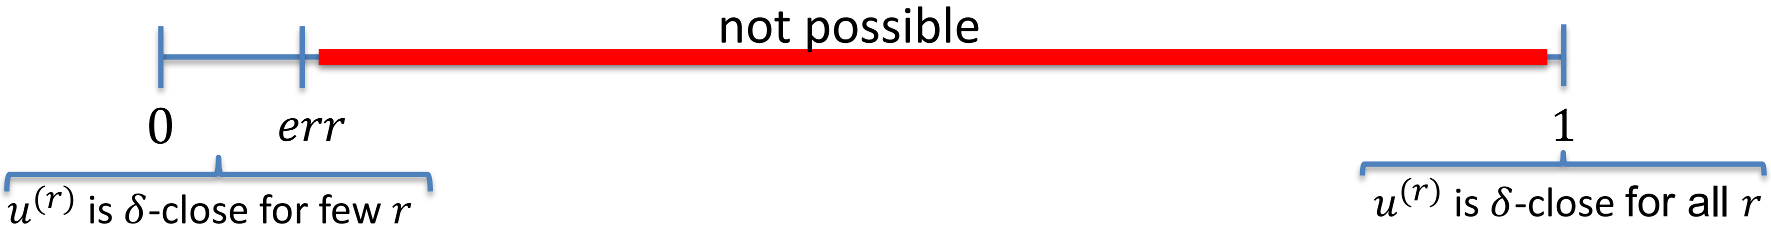

Proximity gaps for other linear codes?

A similar proximity gap holds for every linear code.

min. distance

Thm: (Zeilberger'24) Let $\mathcal{C} \subseteq \mathbb{F}^n$ be an $[n, dim, l]_p$ linear code.
Then $\mathcal{C}$ has a correlated proximity gap for $0 < \delta < 1 - \sqrt[4]{\tau}$
and err = $O\left({}^{kn}/_{|\mathbb{F}|}\right)$, where $\tau \coloneqq 1 - (l/n)$.

(For RS-code $\tau \approx \rho$, so this gap is much weaker than BCIKS'20)

This can be used in a $\mathcal{C}$-proximity IOPP (e.g., Basefold, Blaze)

2^{nd} Distance preserving example: 2-way folding

From now on set $\mathcal{L} = \{1, \omega, \omega^2, \dots, \omega^{n-1}\} \subseteq \mathbb{F}$, where

- n is a power of two, and
- ω is an n-th primitive root of unity $(\omega^n = 1)$

  (requires that n divides $|\mathbb{F}| - 1$)

Then:

- $\omega^{n/2} = -1$ so that if $x = \omega^i \in \mathcal{L}$ then $-x = \omega^{i+(n/2)} \in \mathcal{L}$
- $|\mathcal{L}^2| = |\{a^2 : a \in \mathcal{L}\}| = |\mathcal{L}|/2 = n/2$ $\quad (-a,\ a \to a^2)$

2-way folding a polynomial

A folding transformation: let's start with an example.

Let $f(X) = 1 + 2X + 3X^2 + 4X^3 + 5X^4 + 6X^5 \in \mathbb{F}^{<6}[X]$

Define $f_{\text{even}}(X) := 1 + 3X + 5X^2$ and $f_{\text{odd}}(X) := 2 + 4X + 6X^2$

Then: $f(X) = f_{\text{even}}(X^2) + X \cdot f_{\text{odd}}(X^2)$

Define: for $r \in \mathbb{F}$ define $f_{\text{fold},r} := f_{\text{even}} + r \cdot f_{\text{odd}} \in \mathbb{F}^{<3}[X]$

2-way folding a polynomial: more generally

For $f \in \mathbb{F}^{<d}[X]$ (with d even) define:

- $f_{\text{even}}(X^2) := \frac{f(X)+f(-X)}{2}$ and $f_{\text{odd}}(X^2) := \frac{f(X)-f(-X)}{2X}$
- $f_{\text{fold},r}(X) := f_{\text{even}}(X) + r \cdot f_{\text{odd}}(X) \in \mathbb{F}^{<d/2}[X]$

Then: $f(X) = f_{\text{even}}(X^2) + X \cdot f_{\text{odd}}(X^2)$

- for every $a \in \mathbb{F}$: $f_{\text{fold},r}(a^2)$ can be eval given $f(a), f(-a)$
- $\bar{f} \in \text{RS}[\mathbb{F}, \mathcal{L}, d] \Rightarrow \overline{f_{\text{fold},r}} \in \text{RS}[\mathbb{F}, \mathcal{L}^2, d/2]$ ← unchanged rate $= d/|\mathcal{L}|$

Folding an arbitrary word $u: \mathcal{L} \to \mathbb{F}$

For $u: \mathcal{L} \to \mathbb{F}$ and $r \in \mathbb{F}$ define $u_e, u_o, u_{\text{fold},r}: \mathcal{L}^2 \to \mathbb{F}$ as

- for $a \in \mathcal{L}$: $u_e(a^2) := \frac{u(a)+u(-a)}{2}$ and $u_o(a^2) := \frac{u(a)-u(-a)}{2a}$
- for $b \in \mathcal{L}^2$: $u_{\text{fold},r}(b) := u_e(b) + r \cdot u_o(b)$

(recall $|\mathcal{L}^2| = |\mathcal{L}|/2$)

Lemma (distance preservation): for $0 < \delta < 1 - \sqrt{\rho}$

- $u \in \mathrm{RS}[\mathbb{F}, \mathcal{L}, d] \Rightarrow u_{\text{fold},r} \in \mathrm{RS}[\mathbb{F}, \mathcal{L}^2, d/2]$ for all $r \in \mathbb{F}$
- u is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d] \Rightarrow$

$$\Pr_r[\, u_{\text{fold},r} \text{ is } \delta\text{-far from } \mathrm{RS}[\mathbb{F}, \mathcal{L}^2, d/2]\,] \geq 1 - err$$

Folding an arbitrary word $u\colon \mathcal{L} \to \mathbb{F}$

For $u\colon \mathcal{L} \to \mathbb{F}$ and $r \in \mathbb{F}$ define $u_e, u_o, u_{\text{fold},r}\colon \mathcal{L}^2 \to \mathbb{F}$ as

- for $a \in \mathcal{L}$: $u_e(a^2) := \frac{u(a)+u(-a)}{2}$ and $u_o(a^2) := \frac{u(a)-u(-a)}{2a}$
- for $b \in \mathcal{L}^2$: $u_{\text{fold},r}(b) := u_e(b) + r \cdot u_o(b)$

(recall $|\mathcal{L}^2| = |\mathcal{L}|/2$)

Lemma (distance preservation): for $0 < \delta < 1 - \sqrt{\rho}$

- $u \in \mathrm{RS}[\mathbb{F}, \mathcal{L}, d] \Rightarrow u_{\text{fold},r} \in \mathrm{RS}[\mathbb{F}, \mathcal{L}^2, d/2]$ for all $r \in \mathbb{F}$
- $\Pr_r[\, u_{\text{fold},r}$ is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}^2, d/2]\,] > err \Rightarrow$

  u is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$

(contra-positive)

Why is this true?

The first part of the lemma is easy. Let's prove the second part.

- Suppose that $\Pr_r[\, u_{\text{fold},r}$ is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}^2, d/2]\,] > err$
- Then by the BCIKS'20 theorem, there are $g_e, g_o \in \mathrm{RS}[\mathbb{F}, \mathcal{L}^2, d/2]$ that match u_e, u_o on a set $S \subseteq \mathcal{L}^2$ of size $|S| \geq (1-\delta)(n/2)$
- Define $g\colon \mathcal{L} \to \mathbb{F}$ as $g(a) \coloneqq g_e(a^2) + a \cdot g_o(a^2) \in \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$
- Then: $g(a) = u(a)$ for all $a \in \mathcal{L}$ for which $a^2 \in S$ (2$|S|$ values in $\mathcal{L}$)
- But then $\Delta(u, g) \leq 1 - \frac{2|S|}{n} = 1 - \frac{|S|}{n/2} \leq \delta$.

  $\Rightarrow$ u is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$

An important corollary

Let $\mathcal{C} = \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ and $\mathcal{C}' = \mathrm{RS}[\mathbb{F}, \mathcal{L}^2, d/2]$

Corollary: For $u\colon \mathcal{L} \to \mathbb{F}$ (folding does not decrease distance, w.h.p)

- if $\Delta(u, \mathcal{C}) < 1 - \sqrt{\rho}$ then $\Pr_r\left[\, \Delta\left(u_{\mathrm{fold},r}, \mathcal{C}'\right) \geq \Delta(u, \mathcal{C}) \,\right] \geq 1 - err$
- if $\Delta(u, \mathcal{C}) \geq 1 - \sqrt{\rho}$ then $\Pr_r\left[\, \Delta\left(u_{\mathrm{fold},r}, \mathcal{C}'\right) \geq 1 - \sqrt{\rho} \,\right] \geq 1 - err$

Recall: $\Delta(u, \mathcal{C}) \leq \delta \iff u$ is δ-close to $\mathcal{C}$

4-way folding $u: \mathcal{L} \to \mathbb{F}$ (using $i^2 = -1$)

For $u: \mathcal{L} \to \mathbb{F}$ define $u_0, u_1, u_2, u_3 : \mathcal{L}^4 \to \mathbb{F}$ for $a \in \mathcal{L}$ as

$$\begin{pmatrix} 4 \cdot u_0(a^4) \\ 4a \cdot u_1(a^4) \\ 4a^2 \cdot u_2(a^4) \\ 4a^3 \cdot u_3(a^4) \end{pmatrix} := \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & (-i)^2 & (-i)^3 \\ 1 & -1 & 1 & -1 \\ 1 & i & i^2 & i^3 \end{pmatrix} \cdot \begin{pmatrix} u(a) \\ u(ia) \\ u(i^2 a) \\ u(i^3 a) \end{pmatrix}$$

(a degree-4 FFT)

The 4-way fold of u: *for* $r \in \mathbb{F}$ define $u_{4\text{fold},r}: \mathcal{L}^4 \to \mathbb{F}$ as

$$u_{4\text{fold},r}(b) := u_0(b) + r \cdot u_1(b) + r^2 \cdot u_2(b) + r^3 \cdot u_3(b) \quad \text{for } b \in \mathcal{L}^4$$

Evaluating $u_{4\text{fold},r}(X)$ at $b \in \mathcal{L}^4$ requires <u>four</u> evals. of $u(X)$.

4-way folding $u: \mathcal{L} \to \mathbb{F}$ (using $i^2 = -1$)

For $u: \mathcal{L} \to \mathbb{F}$ define $u_0, u_1, u_2, u_3 : \mathcal{L}^4 \to \mathbb{F}$ for $a \in \mathcal{L}$ as

$$\begin{pmatrix} 4 \cdot u_0(a^4) \\ 4a \cdot u_1(a^4) \\ 4a^2 \cdot u_2(a^4) \\ 4a^3 \cdot u_3(a^4) \end{pmatrix} := \begin{pmatrix} 1 & 1 & 1 & 1 \\ 1 & -i & (-i)^2 & (-i)^3 \\ 1 & -1 & 1 & -1 \\ 1 & i & i^2 & i^3 \end{pmatrix} \cdot \begin{pmatrix} u(a) \\ u(ia) \\ u(i^2 a) \\ u(i^3 a) \end{pmatrix}$$

(a degree-4 FFT)

The 4-way fold of *u*: *for* $r \in \mathbb{F}$ define $u_{4\text{fold},r} : \mathcal{L}^4 \to \mathbb{F}$ as

$$u_{4\text{fold},r}(b) \coloneqq u_0(b) + r \cdot u_1(b) + r^2 \cdot u_2(b) + r^3 \cdot u_3(b) \quad \text{for } b \in \mathcal{L}^4$$

<u>**Fact**</u>: the same distance preservation corollary holds for $u_{4\text{fold},r}$

8-way folding $u: \mathcal{L} \to \mathbb{F}$ (using an 8th root of unity)

Can similarly define 8-way folding, or even 2^w folding for $w \geq 3$.

maps $u: \mathcal{L} \to \mathbb{F}$ to $u_{2^w\mathrm{fold},r}: \mathcal{L}^{2^w} \to \mathbb{F}$ ($|\mathcal{L}^{2^w}| = |\mathcal{L}|/2^w$)

(1) evaluating $u_{2^w\mathrm{fold},r}(b)$ requires 2^w evals. of $u(X)$

⇒ uses a degree-2^w FFT (degree-8 FFT for 8-way folding)

(2) the same distance preservation corollary holds for $u_{2^w\mathrm{fold},r}$

Review: 2-way folding an arbitrary word $u: \mathcal{L} \to \mathbb{F}$

For $u: \mathcal{L} \to \mathbb{F}$ and $r \in \mathbb{F}$ define $u_e, u_o, u_{\text{fold},r}: \mathcal{L}^2 \to \mathbb{F}$ as

- for $a \in \mathcal{L}$: $u_e(a^2) := \frac{u(a)+u(-a)}{2}$ and $u_o(a^2) := \frac{u(a)-u(-a)}{2a}$
- for $b \in \mathcal{L}^2$: $u_{\text{fold},r}(b) := u_e(b) + r \cdot u_o(b)$

Can similarly define 2^w-way folding for $w \geq 1$.

Recall $|\mathcal{L}^2| = |\mathcal{L}|/2$

How FRI works

A Reed-Solomon IOP of Proximity (RS-IOPP)

FRI phase 1: commit phase

Let $\mathcal{C} = \text{RS}[\mathbb{F}, \mathcal{L}, d]$, $u_0 : \mathcal{L} \to \mathbb{F}$, and $\delta \in [0,1]$

Prover $\boldsymbol{P}(\mathcal{C}, u_0, \cdot)$

Verfier $\boldsymbol{V}^{u_0}(\mathcal{C})$

Phase 1: (commit)

sample $r_1 \leftarrow \mathbb{F}$

honest prover:
$u_1 := (u_0)_{\text{fold},r_1}$

$u_1 : \mathcal{L}^2 \to \mathbb{F}$

$|\mathcal{L}^2| = |\mathcal{L}|/2$

sample $r_2 \leftarrow \mathbb{F}$

honest prover:
$u_2 := (u_1)_{\text{fold},r_2}$

$u_2 : \mathcal{L}^4 \to \mathbb{F}$

$|\mathcal{L}^4| = |\mathcal{L}|/4$

FRI phase 1: commit phase

Let $\mathcal{C} = \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$, $u_0 : \mathcal{L} \to \mathbb{F}$, and $\delta \in [0,1]$

Prover $\boldsymbol{P}(\mathcal{C}, u_0, \cdot)$

Verfier $\boldsymbol{V}^{u_0}(\mathcal{C})$

Phase 1: (commit)

sample $r_1 \leftarrow \mathbb{F}$

honest prover:
$u_1 \coloneqq (u_0)_{\mathrm{fold}, r_1}$

$u_1 : \mathcal{L}^2 \to \mathbb{F}$

sample $r_t \leftarrow \mathbb{F}$

honest prover:
$u_t \coloneqq (u_{t-1})_{\mathrm{fold}, r_t}$

$u_t : \mathcal{L}^{2^t} \to \mathbb{F}$

$|\mathcal{L}^{2^t}| = |\mathcal{L}|/2^t$

FRI phase 2: query phase

Let $\mathcal{C} = \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$, $u_0 : \mathcal{L} \to \mathbb{F}$, and $\delta \in [0,1]$

Phase 2: (query)

Verfier $\boldsymbol{V}^{u_0,\ldots,u_t}(\mathcal{C}, r_1, \ldots, r_t, u_t)$

$u_1 : \mathcal{L}^2 \to \mathbb{F}$

$u_2 : \mathcal{L}^4 \to \mathbb{F}$

$\vdots$

$u_t : \mathcal{L}^{2^t} \to \mathbb{F}$

for $i = 1,2,\ldots,t$:

spot check that $u_i = (u_{i-1})_{\mathrm{fold}, r_i}$

output yes if $u_t \in \mathrm{RS}[\mathbb{F}, \mathcal{L}^{2^t}, d/2^t]$

[Prover sent Merkle commits to $u_1, \ldots, u_t$]

Note: prover sends short u_t to verifier explicitly

(FRI terminates when u_t is "short enough")

FRI phase 2: query phase

Let $\mathcal{C} = \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$, $u_0 : \mathcal{L} \to \mathbb{F}$, and $\delta \in [0,1]$

Phase 2: (query)

$u_1 : \mathcal{L}^2 \to \mathbb{F}$

$u_2 : \mathcal{L}^4 \to \mathbb{F}$

⋮

$u_t : \mathcal{L}^{2^t} \to \mathbb{F}$

Verfier $\boldsymbol{V}^{u_0,\dots,u_t}(\mathcal{C}, r_1, \dots, r_t, u_t)$

for $i = 1,2,\dots,t$:

spot check that $u_i = (u_{i-1})_{\mathrm{fold},r_i}$

output yes if $u_t \in \mathrm{RS}[\mathbb{F}, \mathcal{L}^{2^t}, {}^{d}/_{2^t}]$

Why is this δ-sound? Intuition: u_0 is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$ $\Rightarrow$

u_1 is "far" from $\mathrm{RS}[\mathbb{F}, \mathcal{L}^2, d/2]$ $\Rightarrow$ … $\Rightarrow$ u_t is "far" from $\mathrm{RS}[\mathbb{F}, \mathcal{L}^{2^t}, d/2^t]$

How to spot check: method 1

Let $\mathcal{C} = \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$, $u_0 : \mathcal{L} \to \mathbb{F}$, and $\delta \in [0,1]$

Phase 2: (query)

$u_1 : \mathcal{L}^2 \to \mathbb{F}$

$u_2 : \mathcal{L}^4 \to \mathbb{F}$

⋮

$u_t : \mathcal{L}^{2^t} \to \mathbb{F}$

$$z = \frac{u_{i-1}(s) + u_{i-1}(-s)}{2} + \mathrm{r_i} \cdot \frac{u_{i-1}(s) - u_{i-1}(-s)}{2s}$$

Verfier $\boldsymbol{V}^{u_0,\ldots,u_t}(\mathcal{C}, r_1, \ldots, r_t, u_t)$

How to check that $u_i = (u_{i-1})_{\mathrm{fold},r_i}$:

Repeat m times:

- choose random $s \in \mathcal{L}^{2^{i-1}}$
- query $u_{i-1}(s)$, $u_{i-1}(-s)$, $u_i(s^2)$
- compute $z \coloneqq (u_{i-1})_{\mathrm{fold},r_i}(s^2)$
- reject if $z \neq u_i(s^2)$

How to spot check in a picture

$m = 2$: (spot check at two random spots per oracle)

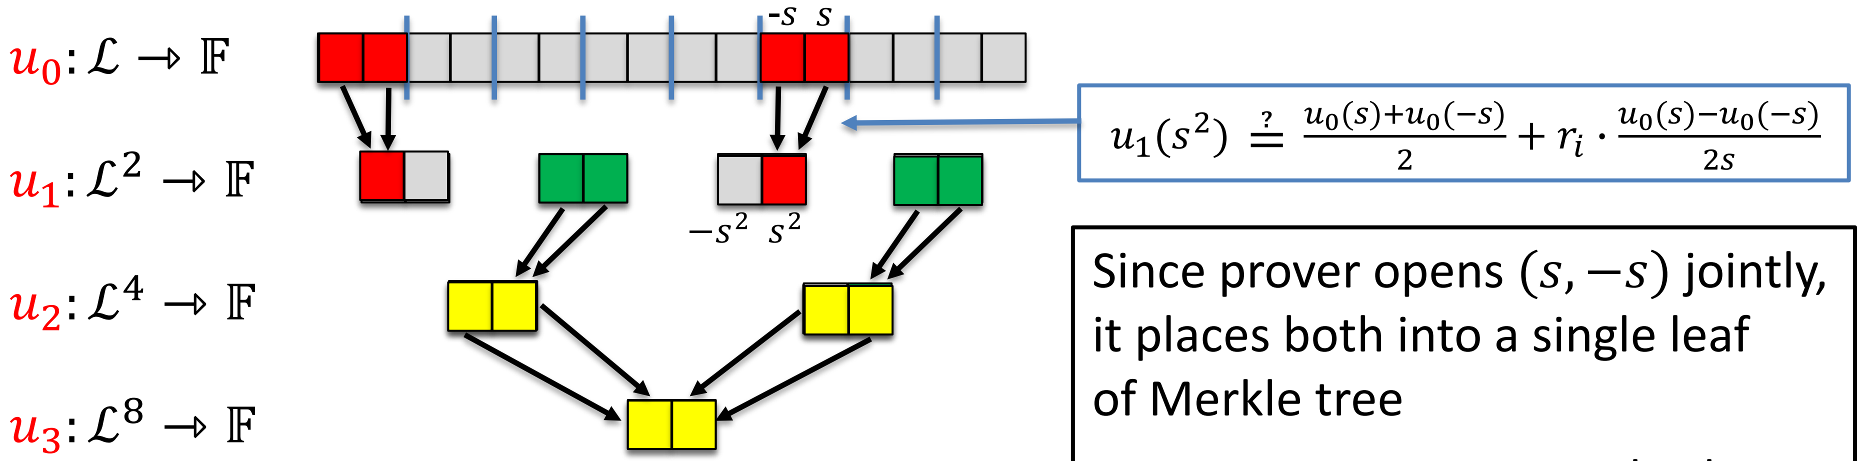

Since prover opens $(s, -s)$ jointly, it places both into a single leaf of Merkle tree

$\Rightarrow$ one query to get both

Reject if any spot checks fail

Total of $2mt$ queries to oracles: $2m$ per inner oracle (u_1, u_2).

Why is this protocol sound?

For $i = 0, \ldots, t$: let

- $\mathcal{C}_i := \mathrm{RS}[\mathbb{F}, \mathcal{L}^{2^i}, d/2^i]$ and
- $\eta_i :=$ (distance of u_i to $\mathcal{C}_i$) $= \Delta(u_i, \mathcal{C}_i) = \min_{w \in \mathcal{C}_i}(\Delta(u_i, w))$

(simplified bound)

Thm: if $0 < \eta_0 < 1 - \sqrt{\rho}$ then Pr[Verifier accepts u_0] ≤ $\left(1 - \frac{1}{2}\eta_0\right)^m$

if $\eta_0 \geq 1 - \sqrt{\rho}$ then Pr[Verifier accepts u_0] ≤ $\left(1 - \frac{1}{2}(1 - \sqrt{\rho})\right)^m$

(recall: m is the number of spot checks per round, and $\rho :=$ (rate of $\mathcal{C}_i$) $= d/|\mathcal{L}|$)

Why is this protocol sound?

Proof idea: To simplify, let's assume that $m = 1$, $\eta_0 < 1 - \sqrt{\rho}$, and

(∗) for all $i = 1, \dots, t$ and $r_i \in \mathbb{F}$: $\Delta((u_{i-1})_{\mathrm{fold},r_i}, \mathcal{C}_i) \geq \Delta(u_{i-1}, \mathcal{C}_{i-1})$

folding does not decrease distance

(note: this only holds w.h.p over $r_i \in \mathbb{F}$ by folding corollary)

Then: $\Pr[\text{accept}] = \prod_{i=1}^{t} \Pr[\text{not reject in round } i] = \prod_{i=1}^{t} [1 - \Delta(u_i, (u_{i-1})_{\mathrm{fold},r_i})]$

independent spot checks per round

prob. u_i is accepted after one spot check

$$\leq \exp(-\textstyle\sum_{i=1}^{t} \Delta(u_i, (u_{i-1})_{\mathrm{fold},r_i})) \leq \exp(-\textstyle\sum_{i=1}^{t} [\Delta((u_{i-1})_{\mathrm{fold},r_i}, \mathcal{C}_i) - \Delta(u_i, \mathcal{C}_i)])$$

$\forall x$: $1 - x \leq e^{-x} = \exp(-x)$

triangular inequality

$\eta_t = 0$ otherwise, Verifier rejects

$$\underset{(*)}{\leq} \exp(-\textstyle\sum_{i=1}^{t} [\Delta(u_{i-1}, \mathcal{C}_{i-1}) - \Delta(u_i, \mathcal{C}_i)]) = \exp(\eta_t - \eta_0) = \exp(-\eta_0) \leq 1 - \eta_0/2$$

How to spot check: method 2 (the FRI method)

Correlated spot checks: spot check starting at a random $s \in \mathcal{L}$

$u_0 : \mathcal{L} \to \mathbb{F}$

$u_1 : \mathcal{L}^2 \to \mathbb{F}$

$u_2 : \mathcal{L}^4 \to \mathbb{F}$

$u_3 : \mathcal{L}^8 \to \mathbb{F}$

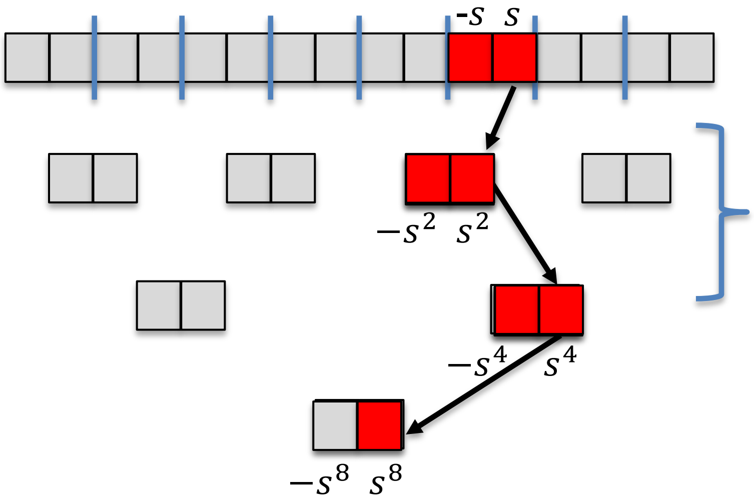

already have $u_1(s^2), u_1(-s^2)$

$\Rightarrow$ only query at $u_2(s^4), u_2(-s^4)$

$\Rightarrow$ 1 query per check

Total of only mt queries to oracles: m per oracle

(recall: method 1 required $2m$ queries per oracle)

How to spot check: the FRI method

Let $\mathcal{C} = \mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$, $u_0 : \mathcal{L} \to \mathbb{F}$

How to check that $u_i = (u_{i-1})_{\text{fold},r_i}$:

Repeat m times:

- choose random $s \in \mathcal{L}$, query $(z_0, y_0) \leftarrow (u_0(s), u_0(-s))$
- for $i = 1, \dots, t$:
  - set $s \leftarrow s^2 \in \mathcal{L}^{2^i}$
  - compute $z := (u_{i-1})_{\text{fold},r_i}(s)$ from z_{i-1}, y_{i-1}
  - query $(z_i, y_i) \leftarrow (u_i(s), u_i(-s))$ ← only one query per round
  - reject if $z \neq z_i$

Why is this sound? (using notation as in earlier proof)

[BKS'18, Thm 7.2]

Proof idea: To simplify, let's assume that $m = 1$, $\eta_0 < 1 - \sqrt{\rho}$, and

(∗) for all $i = 1, \dots, t$ and $r_i \in \mathbb{F}$: $\Delta((u_{i-1})_{\text{fold},r_i}, \mathcal{C}_i) \geq \Delta(u_{i-1}, \mathcal{C}_{i-1})$

folding does not decrease distance

Then: $\Pr[\text{reject}] = \Pr\left[\bigcup_{i=1}^{t}(\text{not reject in round } i)\right] = \sum_{i=1}^{t} \Pr[\text{not reject in round } i]$

can be made into a union of disjoint events

$$= \sum_{i=1}^{t} \Delta(u_i, (u_{i-1})_{\text{fold},r_i}) \geq \sum_{i=1}^{t}\left[\Delta((u_{i-1})_{\text{fold},r_i}, \mathcal{C}_i) - \Delta(u_i, \mathcal{C}_i)\right]$$

triangular inequality

$\eta_t = 0$ otherwise, Verifier rejects

$$\underset{(*)}{\geq} \sum_{i=1}^{t}[\Delta(u_{i-1}, \mathcal{C}_{i-1}) - \Delta(u_i, \mathcal{C}_i)] = \eta_0 - \eta_t = \eta_0 \quad \Rightarrow \quad \Pr[\text{accept}] \leq 1 - \eta_0$$

Comparing methods 1 vs. 2

The FRI method has better soundness: for $\eta_0 = \Delta(u_0, \mathcal{C}_0) < 1 - \sqrt{\rho}$

- **method 1**: Pr[Verifier accepts u_0] ≤ $\left(1 - \frac{1}{2}\eta_0\right)^m$
- **FRI method**: Pr[Verifier accepts u_0] ≤ $(1 - \eta_0)^m$ (lower prob. ⇒ better bound)

To obtain a SNARK via the BCS'16 compiler we need round-by-round soundness

- **method 1**: independent spot checks ⇒ easy to prove R-by-R soundness
- **FRI method**: correlated spot checks ⇒ harder to prove R-by-R soundness

(see [BJKTTZ'23] for R-by-R analysis of FRI)

The number of spot checks m

Goal: u_0 is δ-far from $\mathrm{RS}[\mathbb{F},\mathcal{L},d]$ $\Rightarrow$ Pr[Verifier accepts u_0] $\leq 1/2^{128}$

For $\delta < 1-\sqrt{\rho}$: when u_0 is δ-far from $\mathrm{RS}[\mathbb{F},\mathcal{L},d]$ we know that

Pr[FRI Verifier accepts u_0] $\leq (1-\delta)^m$

So: we want m where $(1-\ \delta)^m \leq 1/2^{128}$

$$\Rightarrow \quad m \geq -128/\log_2(1-\delta)$$

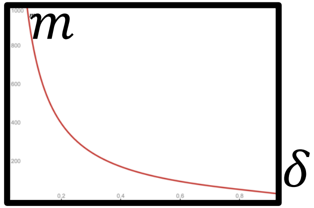

Main point: (1) the bigger δ is, the smaller m needs to be

(2) smaller m $\Rightarrow$ shorter proof and faster verifier

Choosing the code rate $\rho = d/|\mathcal{L}|$

in practice, set $\delta = \delta_{\max} \approx 1 - \sqrt{\rho}$ (to get smallest possible m)

Example 1: $\rho = 1/2$

$\Rightarrow \delta_{\max} \approx 0.29$, $|\mathcal{L}| = 2d$

(Plonky3)

Example 2 : $\rho = 1/4$

$\Rightarrow \delta_{\max} \approx 0.5$, $|\mathcal{L}| = 4d$

Proof length: Longer — Shorter (smaller m)

Prover work: Less — More

shorter codewords $\Rightarrow$ less work for Prover to commit

Is 128-bit security enough??

Suppose m is such that Pr[FRI Verifier accepts a **far** u_0] ≤ 2^{-128}

Fact 1: An adversary that runs FRI 2^{128} times will find a run with favorable spot checks (and forge a proof) with probability $\approx 1/2$

Fact 2: An adversary that runs FRI 2^{80} times will find a run with favorable spot checks (and forge a proof) with probability $\approx 2^{-48}$

For most applications this is sufficient

⇒ do not use less than 120-bits of security; otherwise a 2^{80} adv. will forge proofs.

FRI variants

(1) Higher-order folding
(2) Batch FRI for varying degrees
(3) Reduce proof size by grinding
(4) STIR and WHIR variants

(1) The benefits of higher-order folding

$|\mathcal{L}| = n$

2-way folding **three times**

$u_0 : \mathcal{L} \to \mathbb{F}$

$u_1 : \mathcal{L}^2 \to \mathbb{F}$

$u_2 : \mathcal{L}^4 \to \mathbb{F}$

$u_3 : \mathcal{L}^8 \to \mathbb{F}$

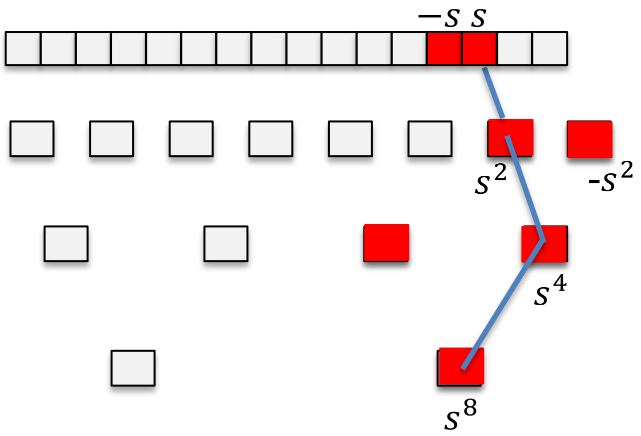

Protocol has $t = \log_2 n$ rounds

But: total of 6 queries per 8-way step

8-way folding **once**

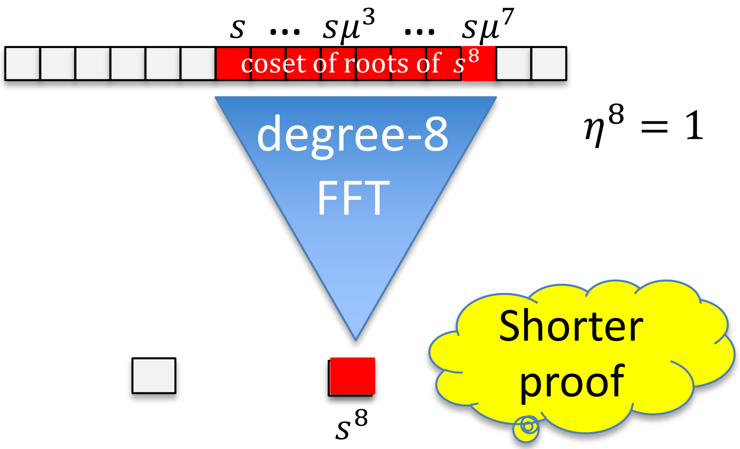

Protocol has $\log_8 n = t/3$ rounds

vs. total of 8 queries per 8-way step

Can shrink Merkle proofs by placing entire coset in one leaf of Merkle tree:

⇒ 3 Merkle proofs are about $3\log_2 n$ hashes | ⇒ 1 Merkle proof is about $\log_2 n$ hashes

(2) Batch FRI for varying degrees: two methods

Poly-IOPP often do multiple RS proximity tests:

Let $u_j : \mathcal{L}_\mathrm{j} \to \mathbb{F}$ be words for $j = 1, \ldots, k$

All $u_1, \ldots, u_k$ are encoded with the same rate $\rho = d_j / |\mathcal{L}_j|$

Goal: reject if for any j, u_j is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}_\mathrm{j}, d_j]$

different d_j

The plan: test all k words as a batch.

Method 1: degree padding

Let $d_{max} := \max_j d_j$, $\mathcal{L}_{\max} := \cup_j \mathcal{L}_j$

$u_1 \colon \mathcal{L}_1 \to \mathbb{F}$  $d_1 = d_{\max}$ $(= \rho \cdot |\mathcal{L}_{\max}|)$

$u_2 \colon \mathcal{L}_2 \to \mathbb{F}$  $d_2 = d_{\max}/4$

$u_3 \colon \mathcal{L}_3 \to \mathbb{F}$  $d_3 = d_{\max}/2$

$u_4 \colon \mathcal{L}_4 \to \mathbb{F}$  $d_4 = d_{\max}$

Honest prover can interpolate all $u_1, \dots, u_k$ to $\mathcal{L}_{\max}$

Method 1: degree padding

Let $d_{max} := \max_j d_j$, $\mathcal{L}_{\max} := \bigcup_j \mathcal{L}_j$

$u_1 \colon \mathcal{L}_{\max} \to \mathbb{F}$ $\quad d_1 = d_{\max} \quad (= \rho \cdot |\mathcal{L}_{\max}|)$

$u_2 \colon \mathcal{L}_{\max} \to \mathbb{F}$ $\quad d_2 = d_{\max}/4$

$u_3 \colon \mathcal{L}_{\max} \to \mathbb{F}$ $\quad d_3 = d_{\max}/2$

$u_4 \colon \mathcal{L}_{\max} \to \mathbb{F}$ $\quad d_4 = d_{\max}$

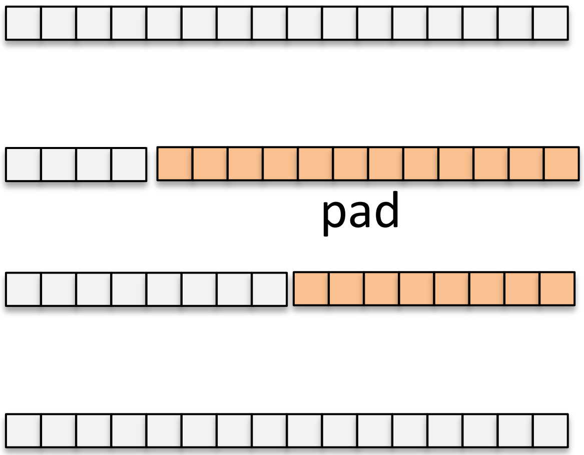

Prover can interpolate all $u_1, \ldots, u_k$ to $\mathcal{L}_{\max}$

Method 1: degree padding

Now, batch as follows: Verifier samples a random $r \in \mathbb{F}$, sends to prover

Honest prover defines: $v_r : \mathcal{L}_{\max} \to \mathbb{F}$ as

$$v_r(a) := \sum_{j=0}^{k} \sum_{i=0}^{d_{max}-d_j} r^{e_{i,j}} \cdot [a^i \cdot u_j(a)]$$

where $e_{i,j}$ is a running counter

Example: suppose $(u_1, d_1), (u_2, d_2), (u_3, d_3)$ s.t. $d_1 = 3$, $d_2 = 4$, $d_3 = 5$.

$$v_r(a) := \underbrace{[u_1(a) + r \cdot au_1(a) + r^2 \cdot a^2 u_1(a)]}_{\text{linear comb. of } u_1(a),\ au_1(a),\ a^2 u_1(a)} + \underbrace{[r^3 \cdot u_2(a) + r^4 \cdot au_2(a)]}_{\text{linear comb. of } u_2(a),\ au_2(a)} + r^5 \cdot u_3(a)$$

linear comb. of $u_1(a)$, $au_1(a)$, $a^2 u_1(a)$ linear comb. of $u_2(a)$, $au_2(a)$

$u_1 \in \mathrm{RS}[\mathbb{F}, \mathcal{L}_1, 3] \Rightarrow a^2 u_1(a) \in \mathrm{RS}[\mathbb{F}, \mathcal{L}_1, 5]$ $u_2 \in \mathrm{RS}[\mathbb{F}, \mathcal{L}_2, 4] \Rightarrow au_2(a) \in \mathrm{RS}[\mathbb{F}, \mathcal{L}_2, 4]$

Method 1: degree padding

Lemma: [STIR, Lemma 4.13] the transform $(u_1, \ldots, u_k; r) \to v_r$ is distance preserving

case 1: (the honest case)

if $\forall j$: $u_j \in \mathrm{RS}[\mathbb{F}, \mathcal{L}_j, d_j]$ then $v_r \in \mathrm{RS}[\mathbb{F}, \mathcal{L}_{\max}, d_{\max}]$ for all r.

case 2: (the dishonest case)

if some u_j is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}_j, d_j]$ then
v_r is δ-far from $\mathrm{RS}[\mathbb{F}, \mathcal{L}_{\max}, d_{\max}]$, w.h.p over r.

The proof follows directly from the RS proximity gap (the BCIKS'20 theorem)

Prover now uses an RS-IOPP to prove that v_r is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}_{\max}, d_{\max}]$

Method 2: pipelining (no padding or interpolation)

Prover $P(\mathcal{C}, (u_1, u_2, u_3), \cdot)$

Verfier $V^{u_1,u_2,u_3}(\mathcal{C})$

Phase 1: (commit)

suppose $d_2 = d_1/2$ and $d_3 = d_1/4$

sample $r_1 \leftarrow \mathbb{F}$

honest prover:

$w_1 := (u_1)_{\text{fold},r_1} + r_1^2 u_2$

(fold u_2 into w_1)

$w_1 : \mathcal{L}^2 \to \mathbb{F}$

$|\mathcal{L}^2| = |\mathcal{L}|/2$

sample $r_2 \leftarrow \mathbb{F}$

honest prover:

$w_2 := (w_1)_{\text{fold},r_2} + r_2^2 u_3$

(fold u_3 into w_2)

$w_2 : \mathcal{L}^4 \to \mathbb{F}$

$|\mathcal{L}^4| = |\mathcal{L}|/4$

⋮

⋮

(3) Grinding to reduce # of spot checks

Prover $\boldsymbol{P}((\mathcal{C},\delta), u_0, \cdot)$ Verfier $\boldsymbol{V}^{u_0}((\mathcal{C},\delta))$

Commit phase:

$u_1 : \mathcal{L}^2 \to \mathbb{F}$

$u_2 : \mathcal{L}^4 \to \mathbb{F}$

$u_3 : \mathcal{L}^8 \to \mathbb{F}$

MerkleCommits (including u_0)

Query phase: Derive queries for spot checks by hashing all MerkleCommits (Fiat-Shamir)

Goal: reduce the number of spot checks m (to reduce proof size)

The problem: reducing m below the computed bound enables adversary to try multiple MerkleCommits, until it finds a favorable set of spot checks

Grinding to reduce # of spot checks

One option: add a grinding phase after commit phase (often used in FRI)

Commit phase:

$u_1 : \mathcal{L}^2 \to \mathbb{F}$

$u_2 : \mathcal{L}^4 \to \mathbb{F}$

$u_3 : \mathcal{L}^8 \to \mathbb{F}$

MerkleCommits (including u_0)

Grind: Find G s.t. MSB(SHA3(G,MerkleCommits,nonce)) $= 0^{64}$

Query phase: Derive queries for spot checks by hashing **all MerkleCommits AND G**

(Fiat-Shamir)

Nonce prevents adversary from pre-computing G (e.g, nonce = head of blockchain)

Why does grinding help?

Adversary: wants δ-far u_0 for which it can generate an RS-proximity proof

FRI without grinding:

m is set so that $\mathrm{E}[\text{time to find } \delta\text{-far } u_0 \text{ with favorable queries}] \geq 2^{128}$

$\Rightarrow$ time to find a false proximity proof is $\approx 2^{128}$

FRI with grinding: every u_0 attempt takes time $\approx 2^{64}$ to find G

$\Rightarrow$ suffice that $\mathrm{E}[\text{time to find } \delta\text{-far } u_0 \text{ with favorable queries}] \geq 2^{64}$

$\Rightarrow$ can halve the number of spot checks m $\left({}^{128}/_{-\log_2(1-\delta)} \rightarrow {}^{64}/_{-\log_2(1-\delta)} \right)$

$\Rightarrow$ shrink proof length by about $\times 2$

(4) STIR: an FRI variant [ACFY'24]

Recall: in FRI, distances and # spot checks m are fixed round-to-round

input: u_0

distance u_0 to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$:

round 1: (4-way fold)

distance u_1 to $\mathrm{RS}[\mathbb{F}, \mathcal{L}^4, d/4]$:

round 2: (4-way fold)

distance u_2 to $\mathrm{RS}[\mathbb{F}, \mathcal{L}^{16}, d/16]$:

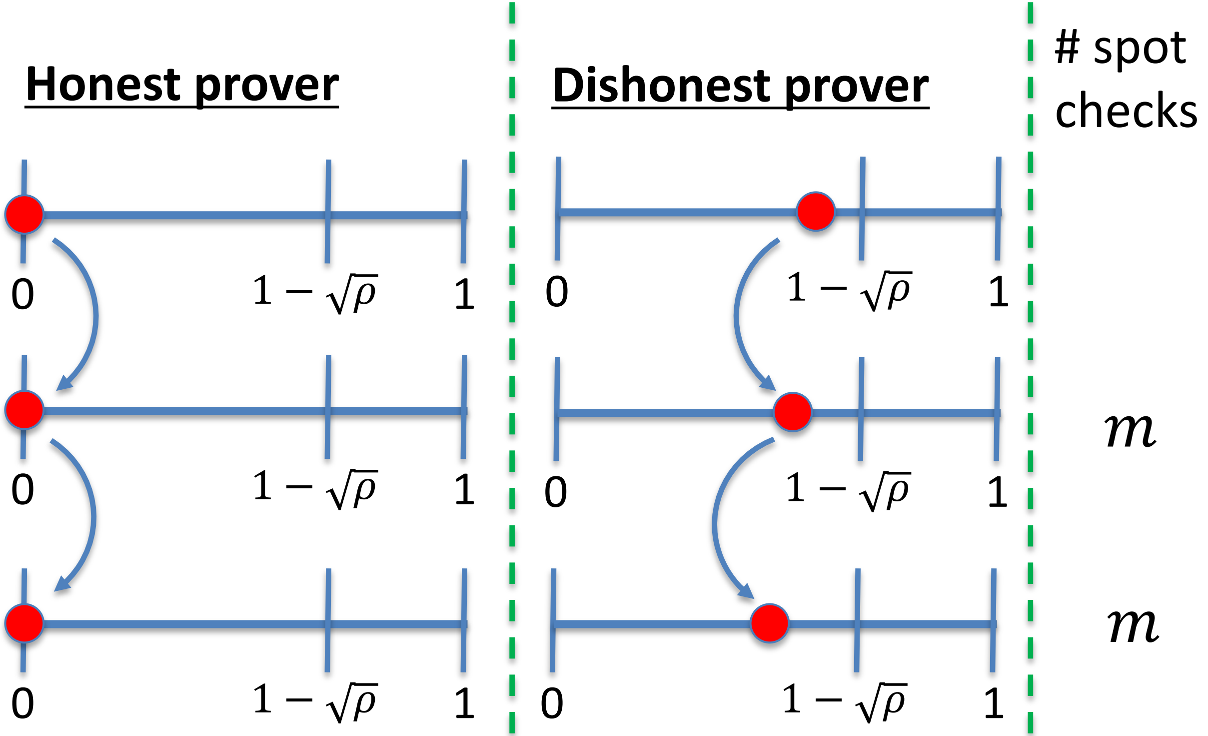

STIR: an FRI variant [ACFY'24]

STIR main idea: in each round reduce the rate and increase distance

⇒ # spot checks can be decreased from round to round

input: u_0

distance u_0 to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$:

(rate ρ)

round 1: (4-way fold)

distance u_1 to $\mathrm{RS}[\mathbb{F}, \mathcal{L}_1, d/4]$:

(rate $\rho/2$)

round 2: (4-way fold)

distance u_2 to $\mathrm{RS}[\mathbb{F}, \mathcal{L}_2, d/16]$:

(rate $\rho/4$)

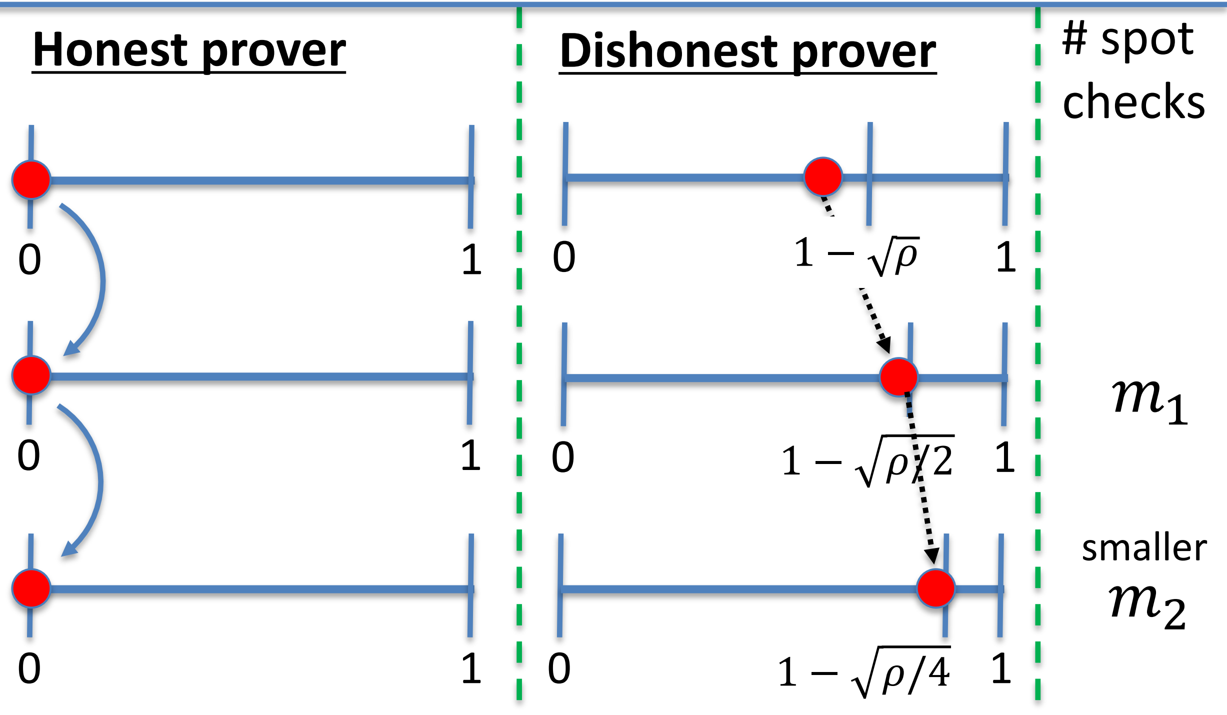

STIR: an FRI variant [ACFY'24]

Idea 1: reduce the code rate by making the honest prover interpolate

Goal: prove u_0 is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$

rate = $\rho = d/|\mathcal{L}|$

$u_0 \colon \mathcal{L} \to \mathbb{F}$

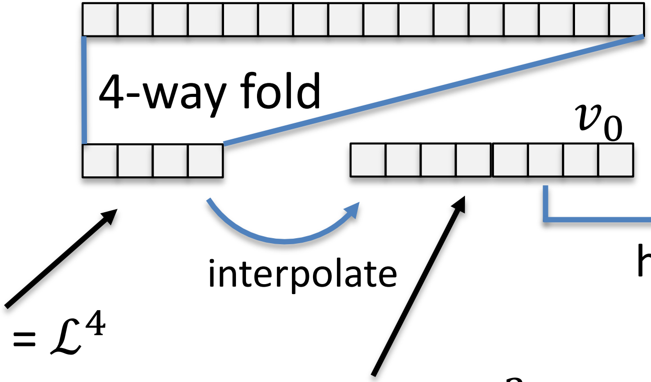

send $v_0 \colon \omega\mathcal{L}^2 \to \mathbb{F}$ to Verifier

honest prover interpolates v_0 from $(u_0)_{4\mathrm{fold},r}$

(half the size of $\mathcal{L}$)

New goal: prove v_0 is δ-close to $\mathrm{RS}[\mathbb{F}, \omega\mathcal{L}^2, d/4]$

rate $= (^{d}/_{4})/(^{|\mathcal{L}|}/_{2}) = \rho/2 \quad \Rightarrow$ lower rate

STIR: an FRI variant [ACFY'24]

Idea 1: reduce the code rate by making the honest prover interpolate

$u_0 : \mathcal{L} \to \mathbb{F}$

Goal: prove u_0 is δ-close to $\mathrm{RS}[\mathbb{F}, \mathcal{L}, d]$

rate = $\rho = d/|\mathcal{L}|$

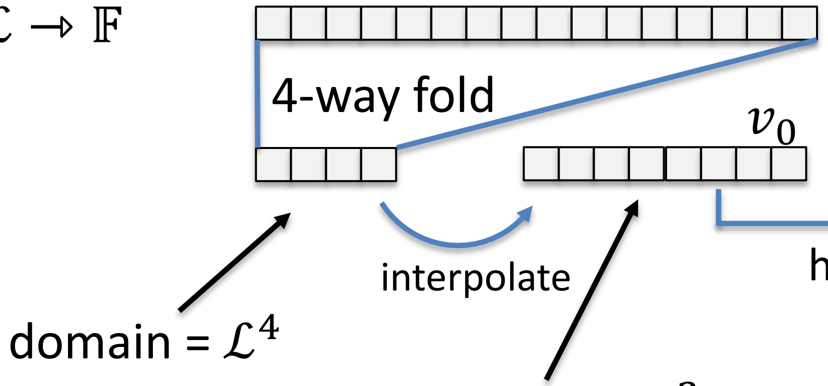

(half the size of $\mathcal{L}$)

send $v_0 : \omega\mathcal{L}^2 \to \mathbb{F}$ to Verifier

honest prover interpolates v_0 from $(u_0)_{4\text{fold},r}$

Main point: Lower rate $\rho/2$ means # spot checks for new goal v_0 can be smaller.

Rate drops by a factor of 2 after every folding step $\Rightarrow$ shorter overall proof

STIR: an FRI variant [ACFY'24]

The problem: now we cannot spot check $v_0: \omega\mathcal{L}^2 \to \mathbb{F}$

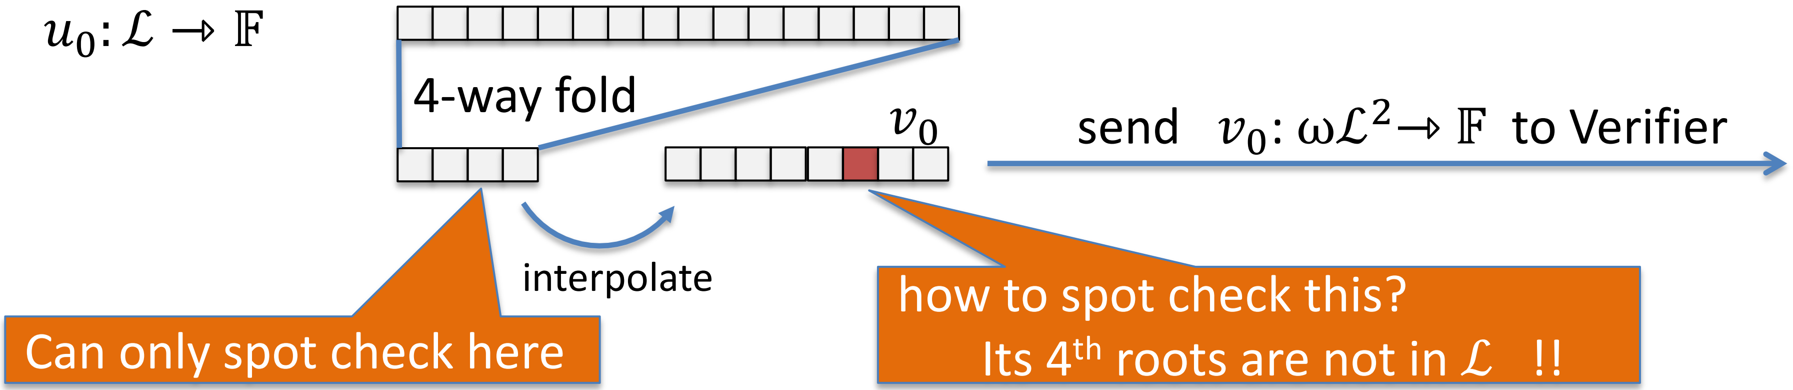

Idea 2: use quotienting for two things:

(1) spot checks on v_0, and

(2) for a malicious prover, increase distance to $\mathrm{RS}[\mathbb{F}, \omega\mathcal{L}^2, d/4]$

How to spot check v_0 by quotienting

How? Honest prover will quotient v_0 by the query points

$u_0 \colon \mathcal{L} \to \mathbb{F}$

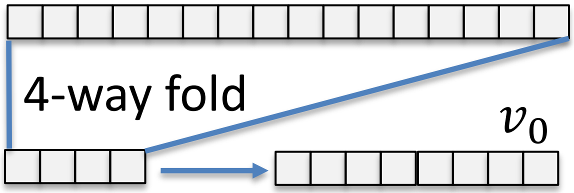

<u>Verfier ***V***</u>

send $v_0 \colon \omega\mathcal{L}^2 \to \mathbb{F}$ to Verifier →

← Verifier sends an out of domain query $t \in \mathbb{F} \setminus \mathcal{L}^4$

Prover sends back $y \coloneqq f(t)$ →

force prover to choose
one $f \in \mathbb{F}^{<d/4}[X]$ s.t.
$\bar{f} \in \mathrm{List}\left[v_0, d/4,\ 1 - \sqrt{\rho/2}\ \right]$

Honest prover will use $f \in \mathbb{F}^{<d/4}[X]$ s.t. $\bar{f} = (u_0)_{\text{4fold},r}$ on $\mathcal{L}^4$

How to spot check v_0 by quotienting

How? Honest prover will quotient v_0 by the query points

$u_0 : \mathcal{L} \to \mathbb{F}$

4-way fold

v_0

<u>Verfier ***V***</u>

send $v_0 : \omega\mathcal{L}^2 \to \mathbb{F}$ to Verifier

Verifier sends an out of domain query $t \in \mathbb{F} \setminus \mathcal{L}^4$

Prover sends back $y := f(t)$

Verifier sends spot check points $s_1, \dots, s_m \in \mathcal{L}^4$

Prover sends back $y_i := f(s_i),\ i = 1, \dots, m$

Honest prover defines u_1 as the result of quotienting v_0 by $\{ (t, y), (s_1, y_1), \dots, (s_m, y_m) \}$

$u_1 : \omega\mathcal{L}^2 \to \mathbb{F}$

How to spot check v_0 by quotienting

How? Honest prover will quotient v_0 by the query points

$u_0 : \mathcal{L} \to \mathbb{F}$

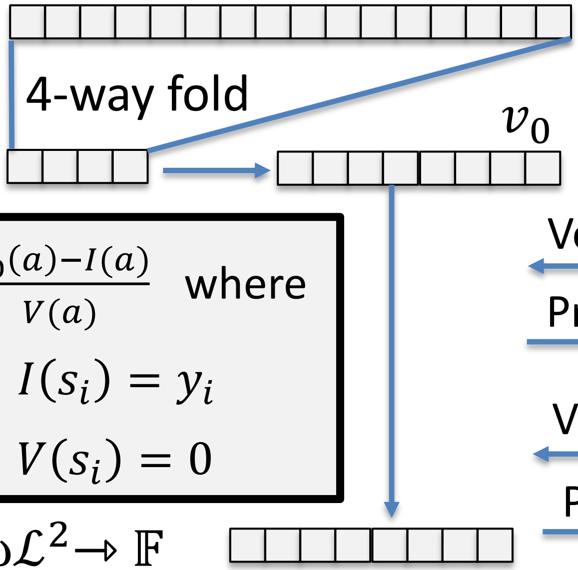

<u>Verfier ***V***</u>

send $v_0 : \omega\mathcal{L}^2 \to \mathbb{F}$ to Verifier →

← Verifier sends an out of domain query $t \in \mathbb{F} \setminus \mathcal{L}^4$

Prover sends back $y \coloneqq f(t)$ →

← Verifier sends spot check points $s_1, \ldots, s_m \in \mathcal{L}^4$

Prover sends back $y_i \coloneqq f(s_i),\ i = 1, \ldots, m$ →

That is, $u_1(\mathrm{a}) \coloneqq \frac{v_0(a) - I(a)}{V(a)}$ where

- $I(t) = y$ and $I(s_i) = y_i$
- $V(t) = 0$ and $V(s_i) = 0$

$u_1 : \omega\mathcal{L}^2 \to \mathbb{F}$

In query phase, Verifier computes y_i itself by querying u_0 at $4m$ points and folding

How to spot check v_0 by quotienting

How? Honest prover will quotient v_0 by the query points

$u_0 \colon \mathcal{L} \to \mathbb{F}$

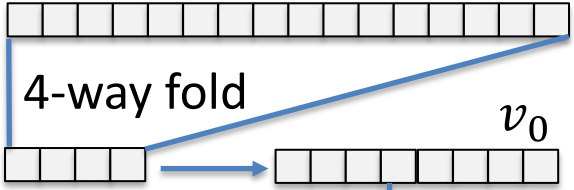

__Verfier *V*__

send $v_0 \colon \omega\mathcal{L}^2 \to \mathbb{F}$ to Verifier

Verifier sends an out of domain query $t \in \mathbb{F} \setminus \mathcal{L}^4$

Prover sends back $y := f(t)$

Verifier sends spot check points $s_1, \ldots, s_m \in \mathcal{L}^4$

Prover sends back $y_i := f(s_i),\ i = 1, \ldots, m$

That is, $u_1(\mathrm{a}) := \frac{v_0(a) - I(a)}{V(a)}$ where

- $I(t) = y$ and $I(s_i) = y_i$
- $V(t) = 0$ and $V(s_i) = 0$

$u_1 \colon \omega\mathcal{L}^2 \to \mathbb{F}$

Iterate to prove that u_1 is $(1 - \sqrt{\rho/2}\,)$-close to $\mathrm{RS}[\mathbb{F}, \omega\mathcal{L}^2, d/4]$ with <u>a smaller</u> m

STIR: summary

Main benefit: STIR proof is about 2× shorter than FRI proof

(using the same rate ρ for the input u_0)

Cons:

- Prover is a bit slower because of interpolation and quotienting
- Verifier is a bit slower because of quotienting
- Batching via pipelining is more cumbersome:
  - Often, functions to batch are all defined using the same rate ρ, but STIR iterations use a different rate in every round

    $\Rightarrow$ Prover will need to interpolate functions to expand to lower rate

WHIR: better than STIR [ACFY'24]

WHIR: combines all spot checks into a Sumcheck $\Rightarrow$ no quotienting

- Fold k levels per round, but Verifier now does fewer field ops.

  $\Rightarrow$ fast verifier (≈1.9M gas in the EVM)

- Supports queries to a multilinear polynomial (not just univariate)

How? Not today. (Builds on BaseFold)

The future: other codes

Are there better codes than Reed-Solomon?

The problem with RS-based SNARKs

- **Not field agnostic**: requires an n-th primitive root of unity in $\mathbb{F}$

  $\Rightarrow$ can only use fields where $n = |\mathcal{L}|$ divides $|\mathbb{F}| - 1$

  $\Rightarrow$ difficult to support specific fields (e.g., for ECDSA arithmetic)

- **Encoding is done via an FFT**: takes time $O(n \cdot \log n)$

  $\Rightarrow$ when $n \approx 2^{20}$, the $(\log_2 n)$ causes $20\times$ work for prover

- **FRI enables:** a (univariate Poly-IOP) ⇾ IOP compiler

  $\Rightarrow$ what about (multilinear Poly-IOP) ⇾ IOP compiler??

  e.g., $g(x_1, x_2, x_3) = 5x_1 + 2x_2 + 4x_1x_2 + 12x_1x_3 + 7x_1x_2x_3$

FRI-like proximity proof for other linear codes

FRI can be generalized to any $[n, k, l]_p$ **linear code** $\mathcal{C} \subseteq \mathbb{F}^n$ where:

There is a sequence of linear codes $\mathcal{C} = \mathcal{C}_0, \ \mathcal{C}_1, \ \ldots, \mathcal{C}_t$ s.t.

1. $\mathcal{C}_\mathrm{i}$ is a $[n_i, k_i, l_i]_p$ linear code where $n = n_0 > n_1 > \cdots > n_t$ and n_t is "sufficiently small",
2. there are distance preserving maps $\mathcal{C}_{i-1} \rightarrow \mathcal{C}_i$ for $i = 1, \ldots, t$,
3. there is a "fast" encoding algorithm $\mathbb{F}^k \rightarrow \mathcal{C}$, and
4. min-distance of $\mathcal{C}_0$ is sufficiently large

(to reduce # of spot checks)

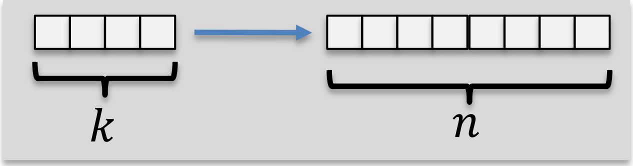

A proximity proof for other linear codes

A **field agnostic** proximity test: (e.g., FRI over the ECDSA prime)

- Gives a (univariate Poly-IOPP) → IOPP over an arbitrary prime p

(1) ECFFT [BCKL'22]:

FRI using functions over an elliptic curve $E/\mathbb{F}_p$,
where the order of $E(\mathbb{F}_p)$ is divisible by n

(even though p is not)

(2) a proximity proof for algebraic geometric codes [BLNR'20].
Here polynomials are replaced with "functions on a curve"

A proximity proof for other linear codes

Circle Stark [HLP'24]: Let $M_{11} := 2^{31} - 1$ (the 11th Mersenne prime)

arithmetic mod M_{11} is super fast, ($x \; mod \; M_{11}$ is just an addition)
but $M_{11} - 1$ is not divisible by a high power of 2

Instead: run FRI over the *projective line* mod M_{11} whose size is $M_{11} + 1$

One way to represent the projective line is as points on a circle:

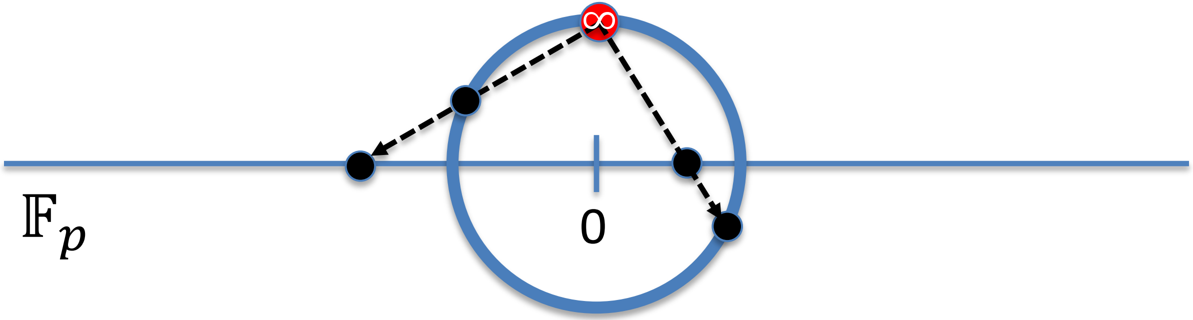

divisible by a high power of 2

So: the circle $(x^2 + y^2 = 1)$ is the same as $\mathbb{F}_p \cup \{\infty\}$ $\Rightarrow$ $(p + 1)$ points

BaseFold [ZCF'23]

BaseFold: generalizes FRI to any **foldable** code

⇒ The generalization gives a **field agnostic** proximity test

[note: every foldable code is a multilinear Reed-Muller code]

For a (multilinear Poly-IOPP) → IOPP compiler need a multilinear PCS

- The problem: quotienting only applies to univariates
- BaseFold solution:

  build a multilinear PCS from Sumcheck and a proximity test

(also adopted into Whir)

How? Not today.

More SNARK-useful linear codes

Spielman codes: [BCG'20, Breakdown'21, Orion'22]

- Linear codes with a good minimum distance and a very fast (linear time) encoding algorithm $\mathbb{F}^k \to \mathcal{C}$.
- Also field-agnostic.

Cons: large IOPP proof $\Rightarrow$ large SNARK proof

Expand Accumulate codes: [BFKTWZ'24]

- Field-agnostic codes, but shorter proofs than Breakdown

Cons: $O(n \cdot \log n)$ time encoding.

More SNARK-useful linear codes

Repeat-Accumulate-Accumulate (RAA) codes: [Blaze'24]

- Constructs a multilinear polynomial commitment over $\mathbb{F}_{2^k}$ with a fast (linear time) prover time and $O(\log^2 n)$ proof size

  $\Rightarrow$ $\mathbb{F}_{2^k}$ is friendly to modern CPU instructions

- The commitment uses the tensor code approach of [BCG'20] (making use of Sumcheck).

Much more to do in non-RS based SNARKs

Further reading

- FRI (2018) and analysis (2018): Fast Reed–Solomon Interactive Oracle Proofs of Proximity
- DEEP-FRI (2019): Out of domain sampling improves soundness
- BCIKS (2020): Proximity Gaps for Reed–Solomon Codes
- CircleSTARK (2024): FRI using a Mersenne prime
- STIR (2024): Reed–Solomon proximity testing with fewer queries
- WHIR (2024): Reed–Solomon proximity testing with a fast verifier

Beyond Reed-Solomon codes (a few recent results):

- Breakdown (2021), Orion (2022): Polynomial commitments with a fast prover
- BaseFold (2023): Efficient Polynomial commitments from foldable codes
- Blaze (2024): Fast SNARKs from Interleaved RAA Codes

END OF MODULE