

STIR



Reed-Solomon Proximity Testing with Fewer Queries

Gal Arnon



Giacomo Fenzi



Alessandro Chiesa



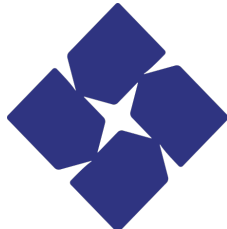
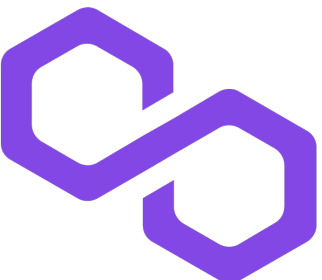

Eylon Yogev



Motivation

STARKs & Friends

- **IOP**-based SNARKs instantiated using the BCS transformation
- Secure in the ROM: can be instantiated with **only** hash-functions
- **Transparent (public-coin) setup**
- Fast proving, not tied to ECC fields, **post quantum secure in QRROM [CMS]**.

Rollups:  **STARKWARE**  **polygon**  **zkSync**

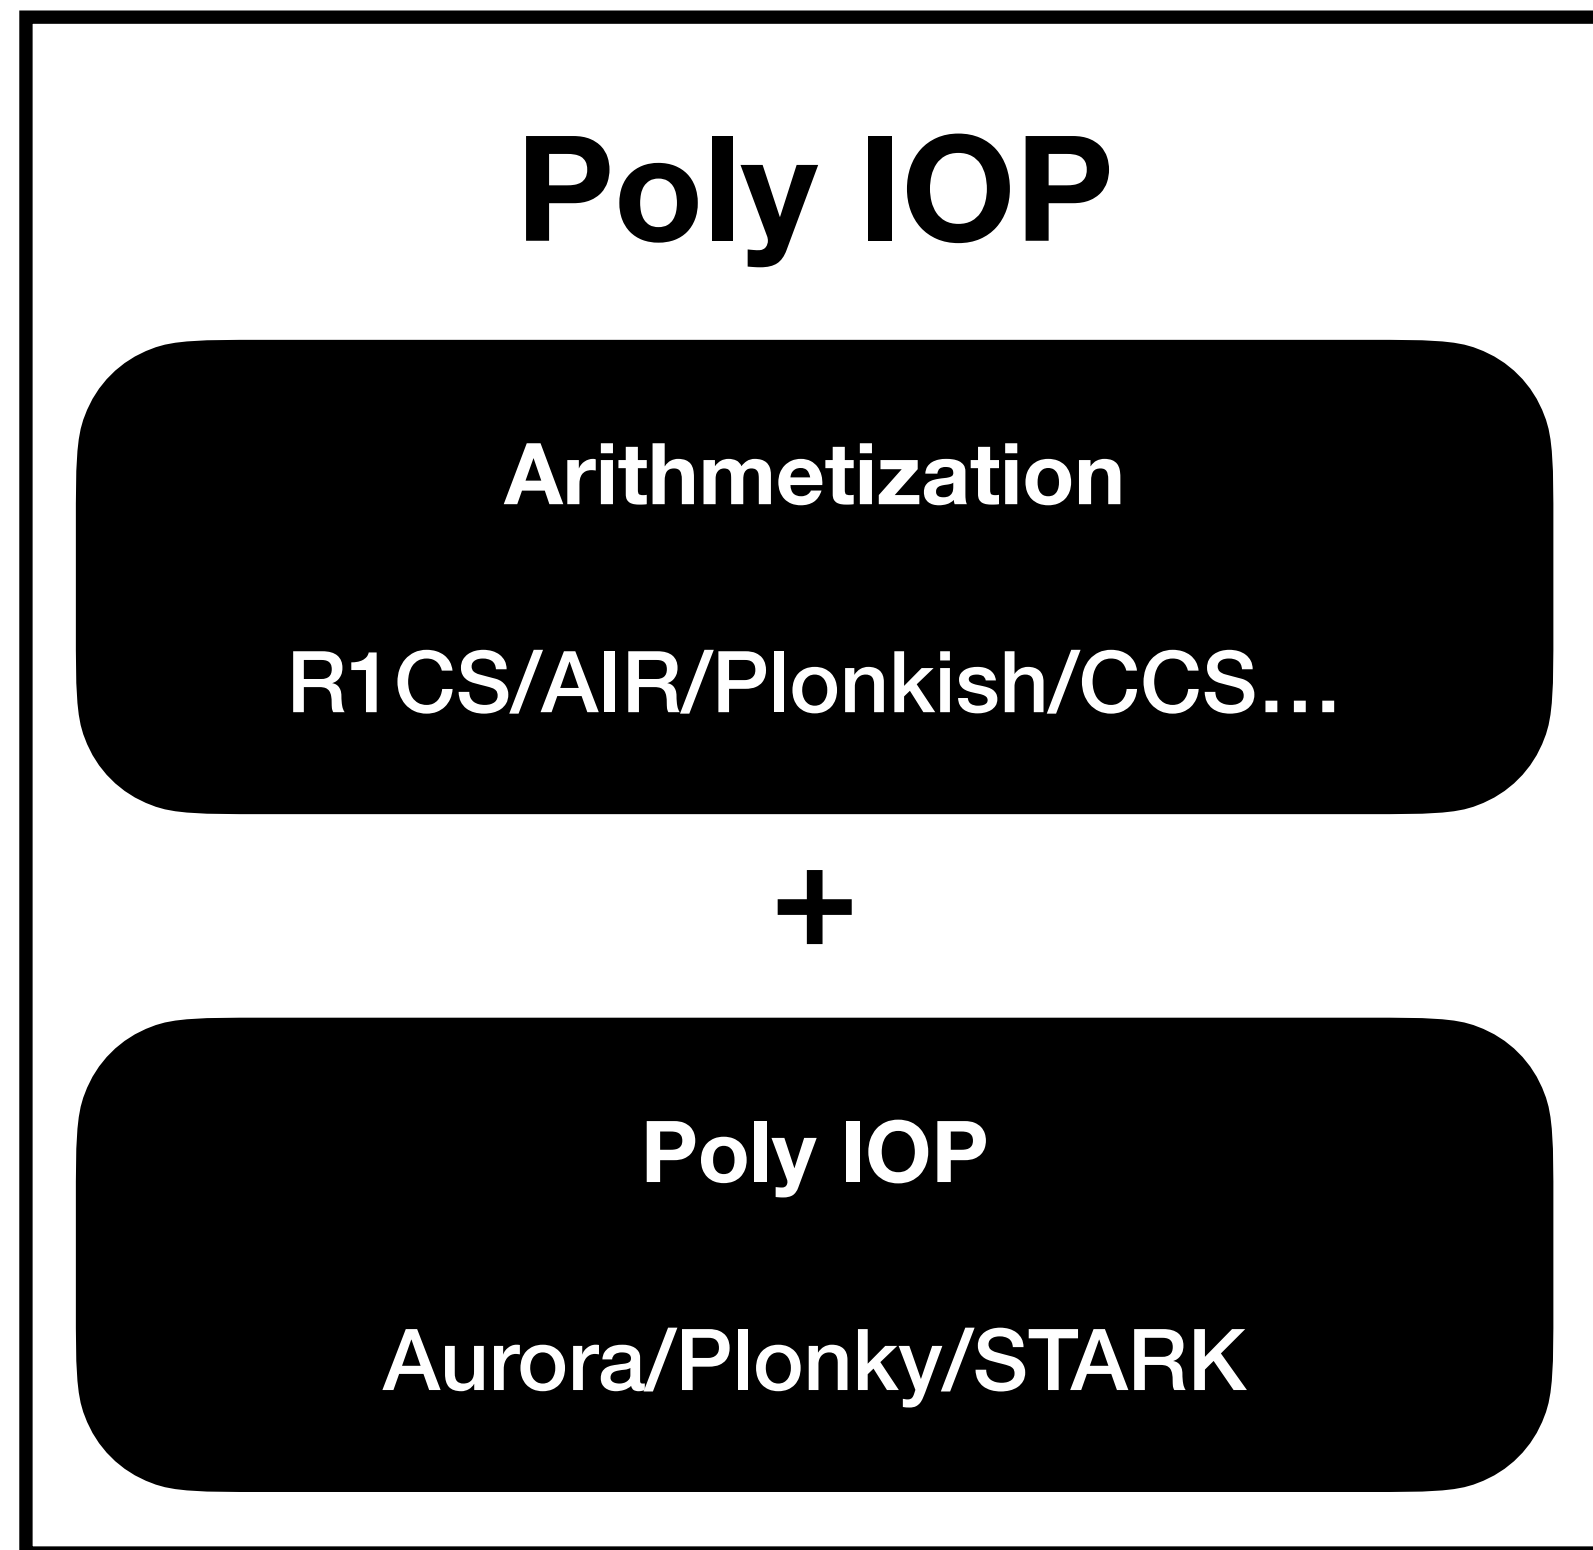
zkVMs:  **Succinct**





And more...

Anatomy of an IOP-based SNARK

Reducing to low-degree testing



Reed-Solomon (RS)
Encoding



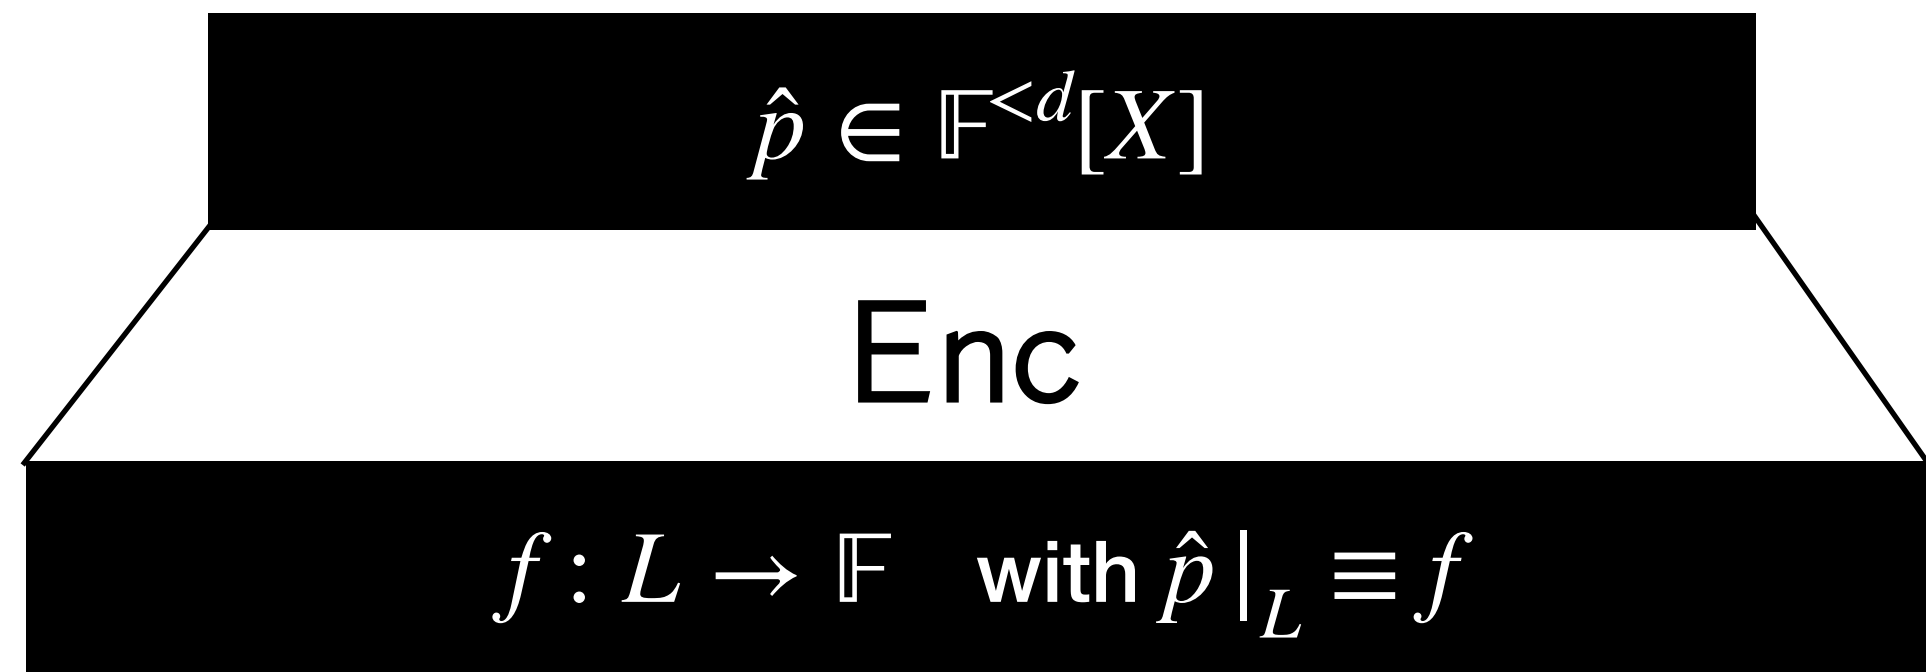

$\gtrsim 80\%$ of the
argument size!

BCS
transformation



Background


RS Codes and IOPPs



Rate: $\rho = d/|L|$, think $\rho = 1/4$

L smooth \implies Enc is an FFT


IOPP for RS



- $f \in RS[\mathbb{F}, L, d] \implies V$ accepts
- f is δ -far from $RS[\mathbb{F}, L, d] \implies V$ rejects w.h.p.

V makes “few queries” to f and proof oracles

FRI Protocol



g_1 is supposed to be the folding of f around α_1

Recursively test the claim that $g_i \in \text{RS}[\mathbb{F}, L^{k^i}, d/k^i]$

Smooth structure allows to check consistency in a single query-phase at the end

Rounds: $O(\log d)$

Proof length: $O(|L|)$

Queries: $O\left(\lambda \cdot \frac{\log d}{-\log \sqrt{\rho}}\right)$ for $\delta = 1 - \sqrt{\rho}$

(To get λ -bits of security, **without conjecture**)

round-by-round

In this talk: $L^k = \{x^k : x \in L\}$

Our results

STIR

An IOPP for RS with

Rounds: $O(\log d)$

Proof length: $O(|L|)$

Queries:

$$O\left(\lambda \cdot \log\left(\frac{\log d}{-\log \sqrt{\rho}}\right) + \log d\right) \quad \text{for } \delta = 1 - \sqrt{\rho}$$

(To get λ -bits of security, **without conjecture**)

round-by-round

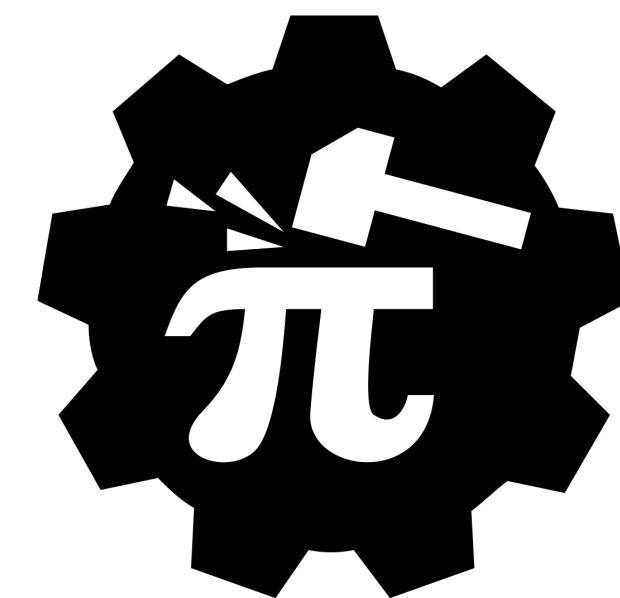
Comparison to FRI

$$\text{FRI: } O\left(\lambda \cdot \frac{\log d}{-\log \sqrt{\rho}}\right)$$

$$\text{STIR: } O\left(\lambda \cdot \log\left(\frac{\log d}{-\log \sqrt{\rho}}\right) + \log d\right)$$

- **Drop-in** replacement of FRI
- **Fewer** queries leads to:
 - Fewer authentication paths \implies **smaller** argument size
 - Smaller verifier hash-complexity \implies **faster** and more **recursion-friendly!**
- Rough query comparison:
 - Example parameters: $\rho = 1/4$, $\delta = 1 - \sqrt{\rho}$, targeting 100-bits of security
 - FRI: \sim **400** queries vs STIR: \sim **200** queries

can increase by PoW



Implementation

- Rust 🦀 implementation, available at [WizardOfMenlo/stir](https://github.com/WizardOfMenlo/stir)
- [Arkworks](#) as backend, 192-bit field for benchmarks, reasonably optimised
- Implemented both FRI and STIR
- Modular over choice of:
 - Field
 - Fiat-Shamir hash
 - Merkle hash
- Decently well written (for academia! 📚)

```
pub trait LowDegreeTest<F, MerkleConfig, FSConfig>
where
    F: FftField,
    MerkleConfig: Config,
    FSConfig: CryptographicSponge,
    FSConfig::Config: Clone,
{
    type Prover: Prover<
        F,
        MerkleConfig,
        FSConfig,
        Commitment = <Self::Verifier as Verifier<F, MerkleConfig, FSConfig>>::Commitment,
        Proof = <Self::Verifier as Verifier<F, MerkleConfig, FSConfig>>::Proof,
    >;
    type Verifier: Verifier<F, MerkleConfig, FSConfig>;


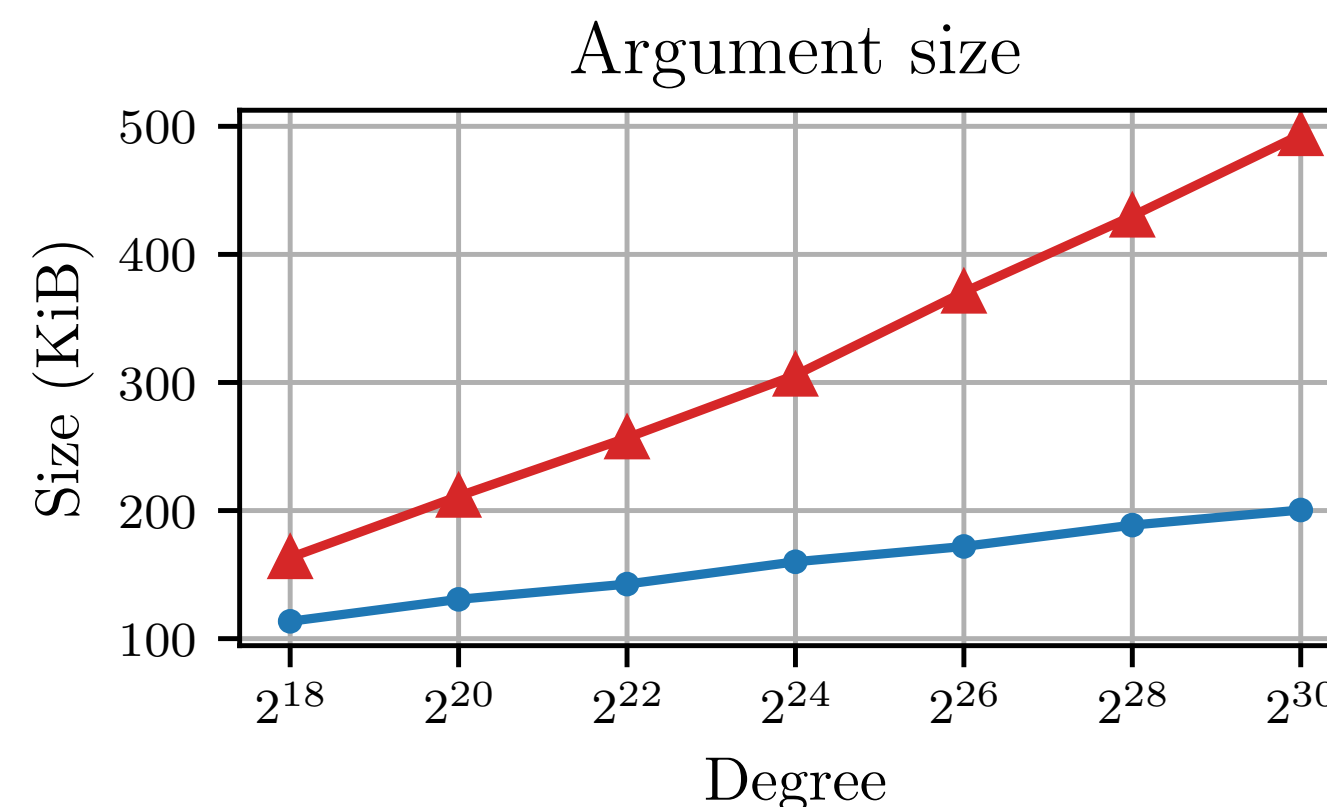
    fn instantiate(
        parameters: Parameters<F, MerkleConfig, FSConfig>,
    ) -> (Self::Prover, Self::Verifier) {
        let prover = Self::Prover::new(parameters.clone());
        let verifier = Self::Verifier::new(parameters);

        (prover, verifier)
    }
}
```

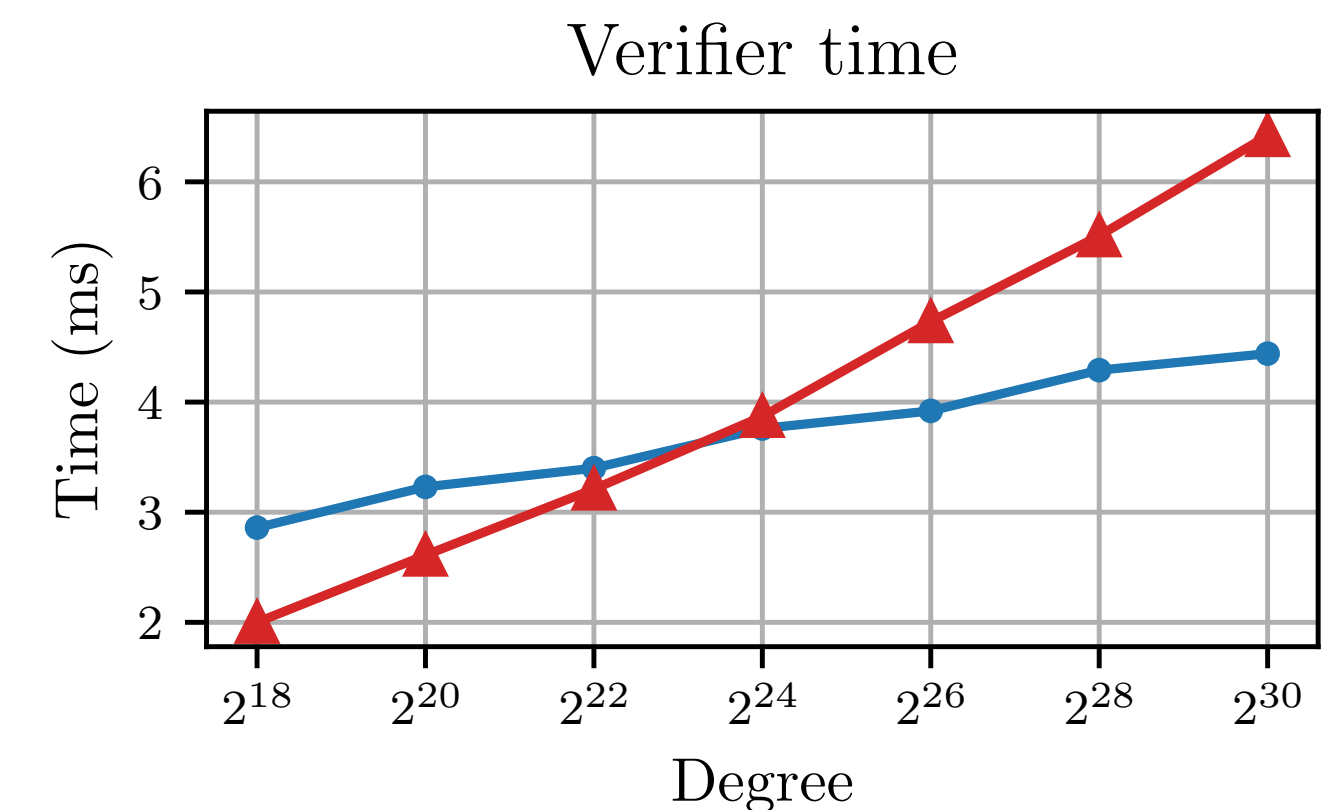
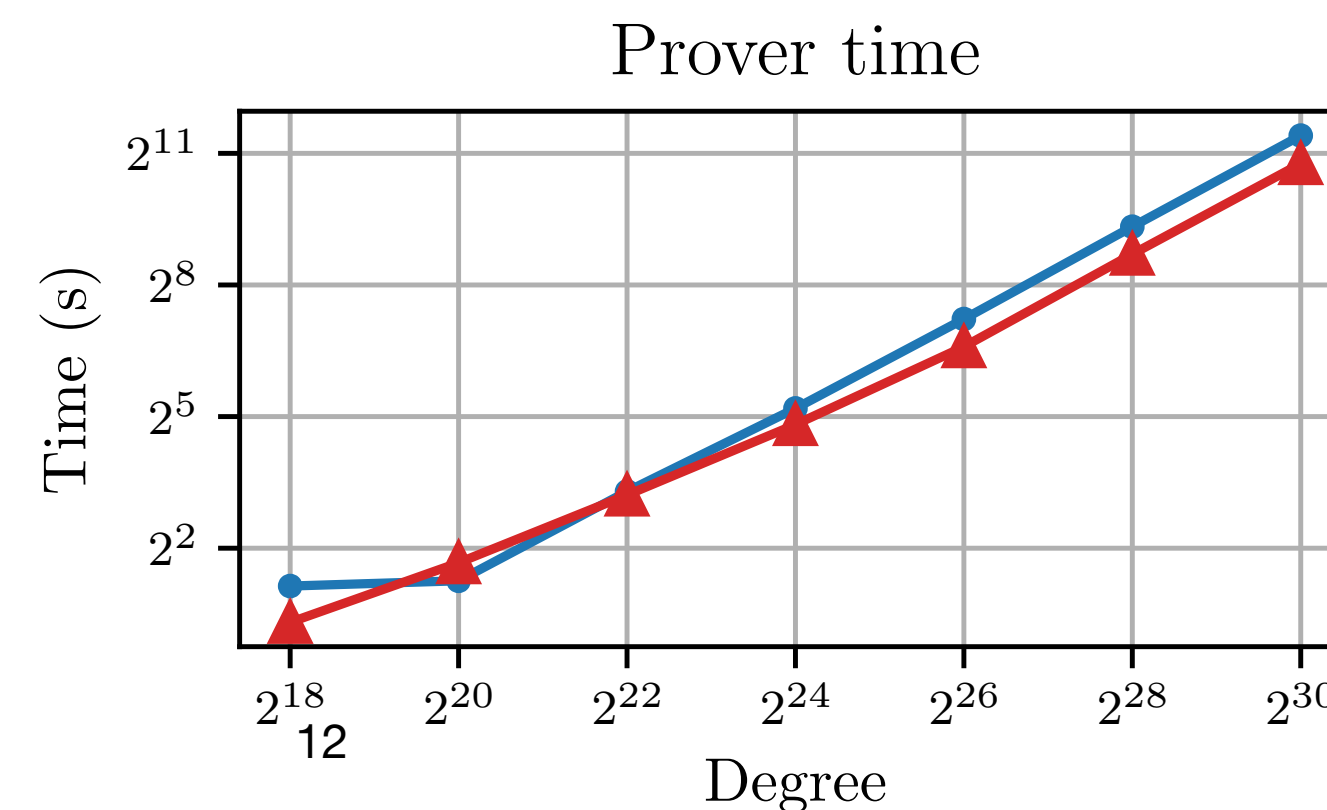
Results

- Compared to FRI: better **argument size** and **verifier hash complexity** across **every** degree-rate pair benchmarked!
- Larger improvements when degree and rate increase

$d = 2^{24}, \rho = 1/4$	FRI	STIR
Size (KiB)	177	107
Hashes	3.5k	1.8k



$d = 2^{30}, \rho = 1/2$	FRI	STIR
Size (KiB)	494	200
Hashes	10k	3.8k



Techniques

Main Idea: Shift-To-Improve-Rate

A smaller rate makes the code easier to test

- A STIR iteration reduces testing proximity to $C = \text{RS}[\mathbb{F}, L, d]$ to testing proximity to $C' = \text{RS}[\mathbb{F}, L', d/k]$ where $|L'| = |L|/2$.
- New rate is $\rho' = (2/k) \cdot \rho$, if $k > 2$ the rate **improves!** (i.e. the new code is easier to test).
- We also **amplify distance**. If f is δ -far then f' is $(1 - \sqrt{\rho'})$ -far to the new code unless with probability $(1 - \delta)^t$ where t is number of queries.

- Round-by-round errors roughly as below so can set:

$$(1 - \delta)^{t_0}, \rho_1^{\frac{t_1}{2}}, \dots, \rho_M^{\frac{t_M}{2}}$$


$$t_0 = \frac{\lambda}{-\log(1 - \delta)}$$

$$t_i = \frac{\lambda}{-\log(\sqrt{\rho_i})}$$

Less queries
each round!

Folding

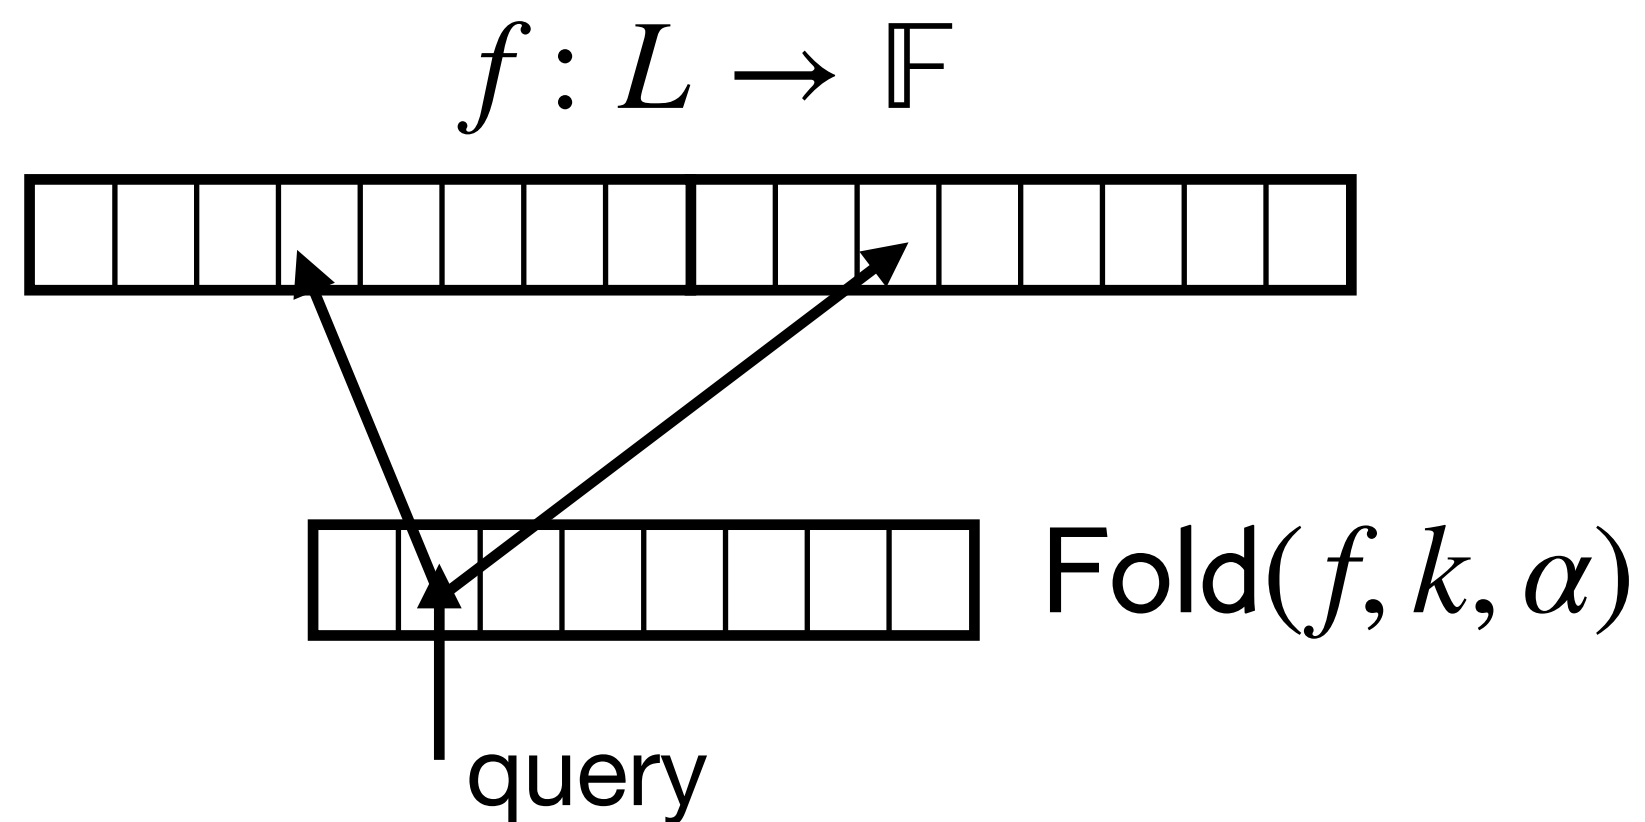
Reduce $RS[\mathbb{F}, L, d]$ to $RS[\mathbb{F}, L^k, d/k]$




Selecting α defines a function $\text{Fold}(f, k, \alpha)$

Local

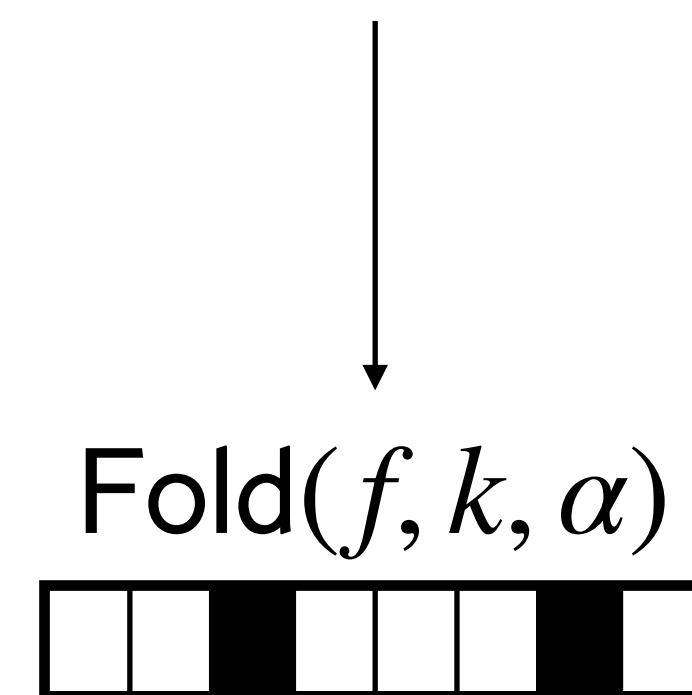
By querying f at k locations, **V** can compute $\text{Fold}(f, k, \alpha)$ at $z \in L^k$



Distance Preserving



f is δ -far from RS



w.h.p. over α

Fold is δ -far from RS

Quotienting

Enforce constraints on f or amplify distance


Let $f : L \rightarrow \mathbb{F}$ be a function
and $\text{Ans} : S \rightarrow \mathbb{F}$ be a list of
(claimed) evaluations of (the
extension of) f on S

$$\text{Quotient}(f, \text{Ans})(x) := \frac{f(x) - \text{Ans}(x)}{V_S(x)}$$

Local

V can compute
 $\text{Quotient}(f, \text{Ans})$ at $x \in L - S$
by querying f at x

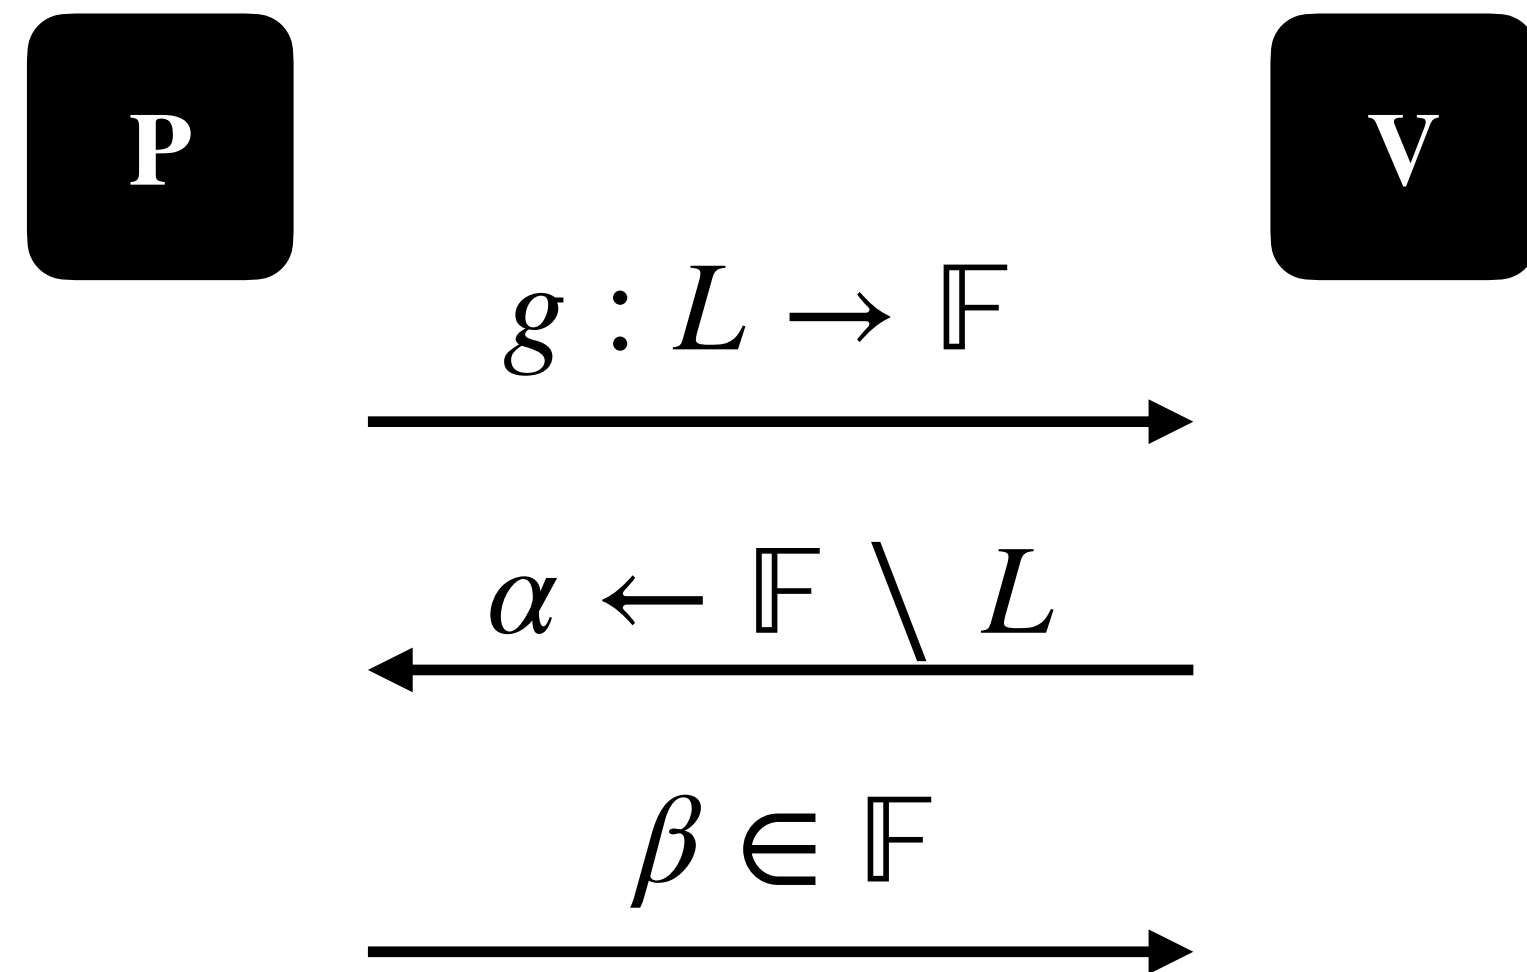
Consistency



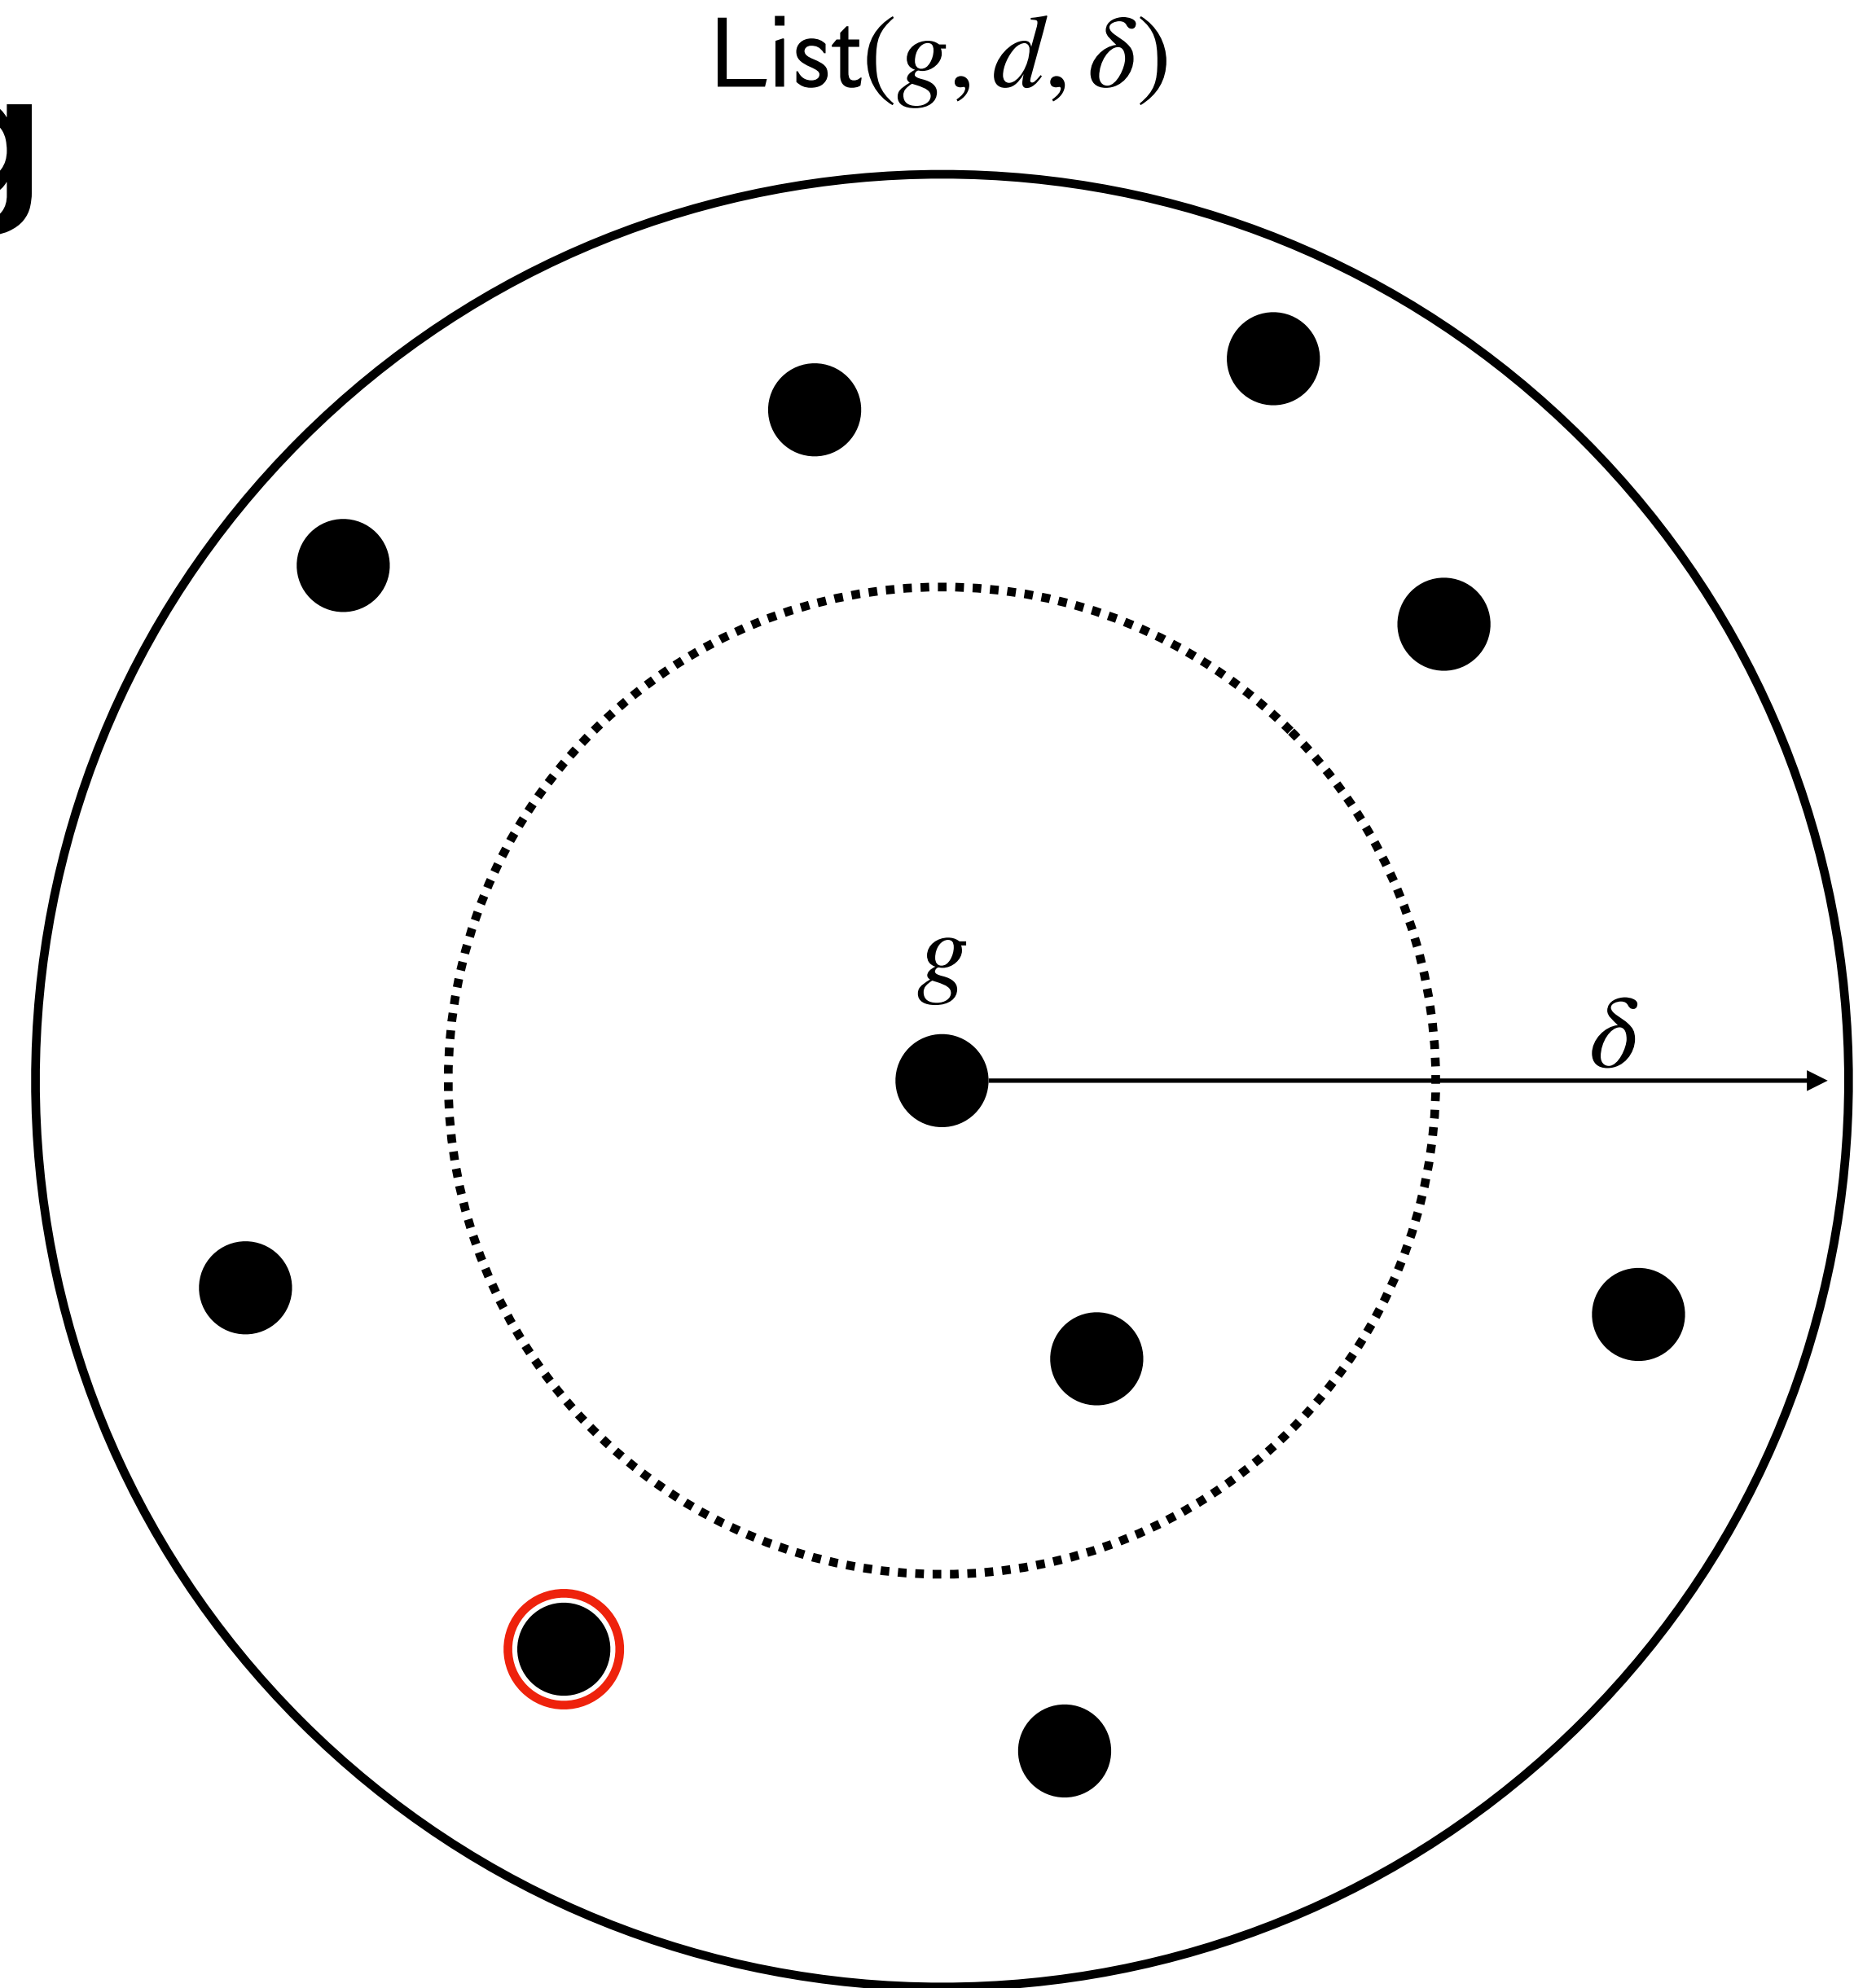
If every $\hat{v} \in \text{List}(f, d, \delta)$
has $\hat{v}|_S \neq \text{Ans}$ then
 $\text{Quotient}(f, \text{Ans})$ is δ -far
from RS

Out Of Domain sampling

Move to unique decoding range



By fundamental theorem of algebra there is at most **one** $\hat{u} \in \text{List}(g, d, \delta)$ such that $\hat{u}(\alpha) = \beta$, w.h.p.



Use $\text{Quotient}(g, \alpha \mapsto \beta)$ to enforce the constraint

STIR iteration

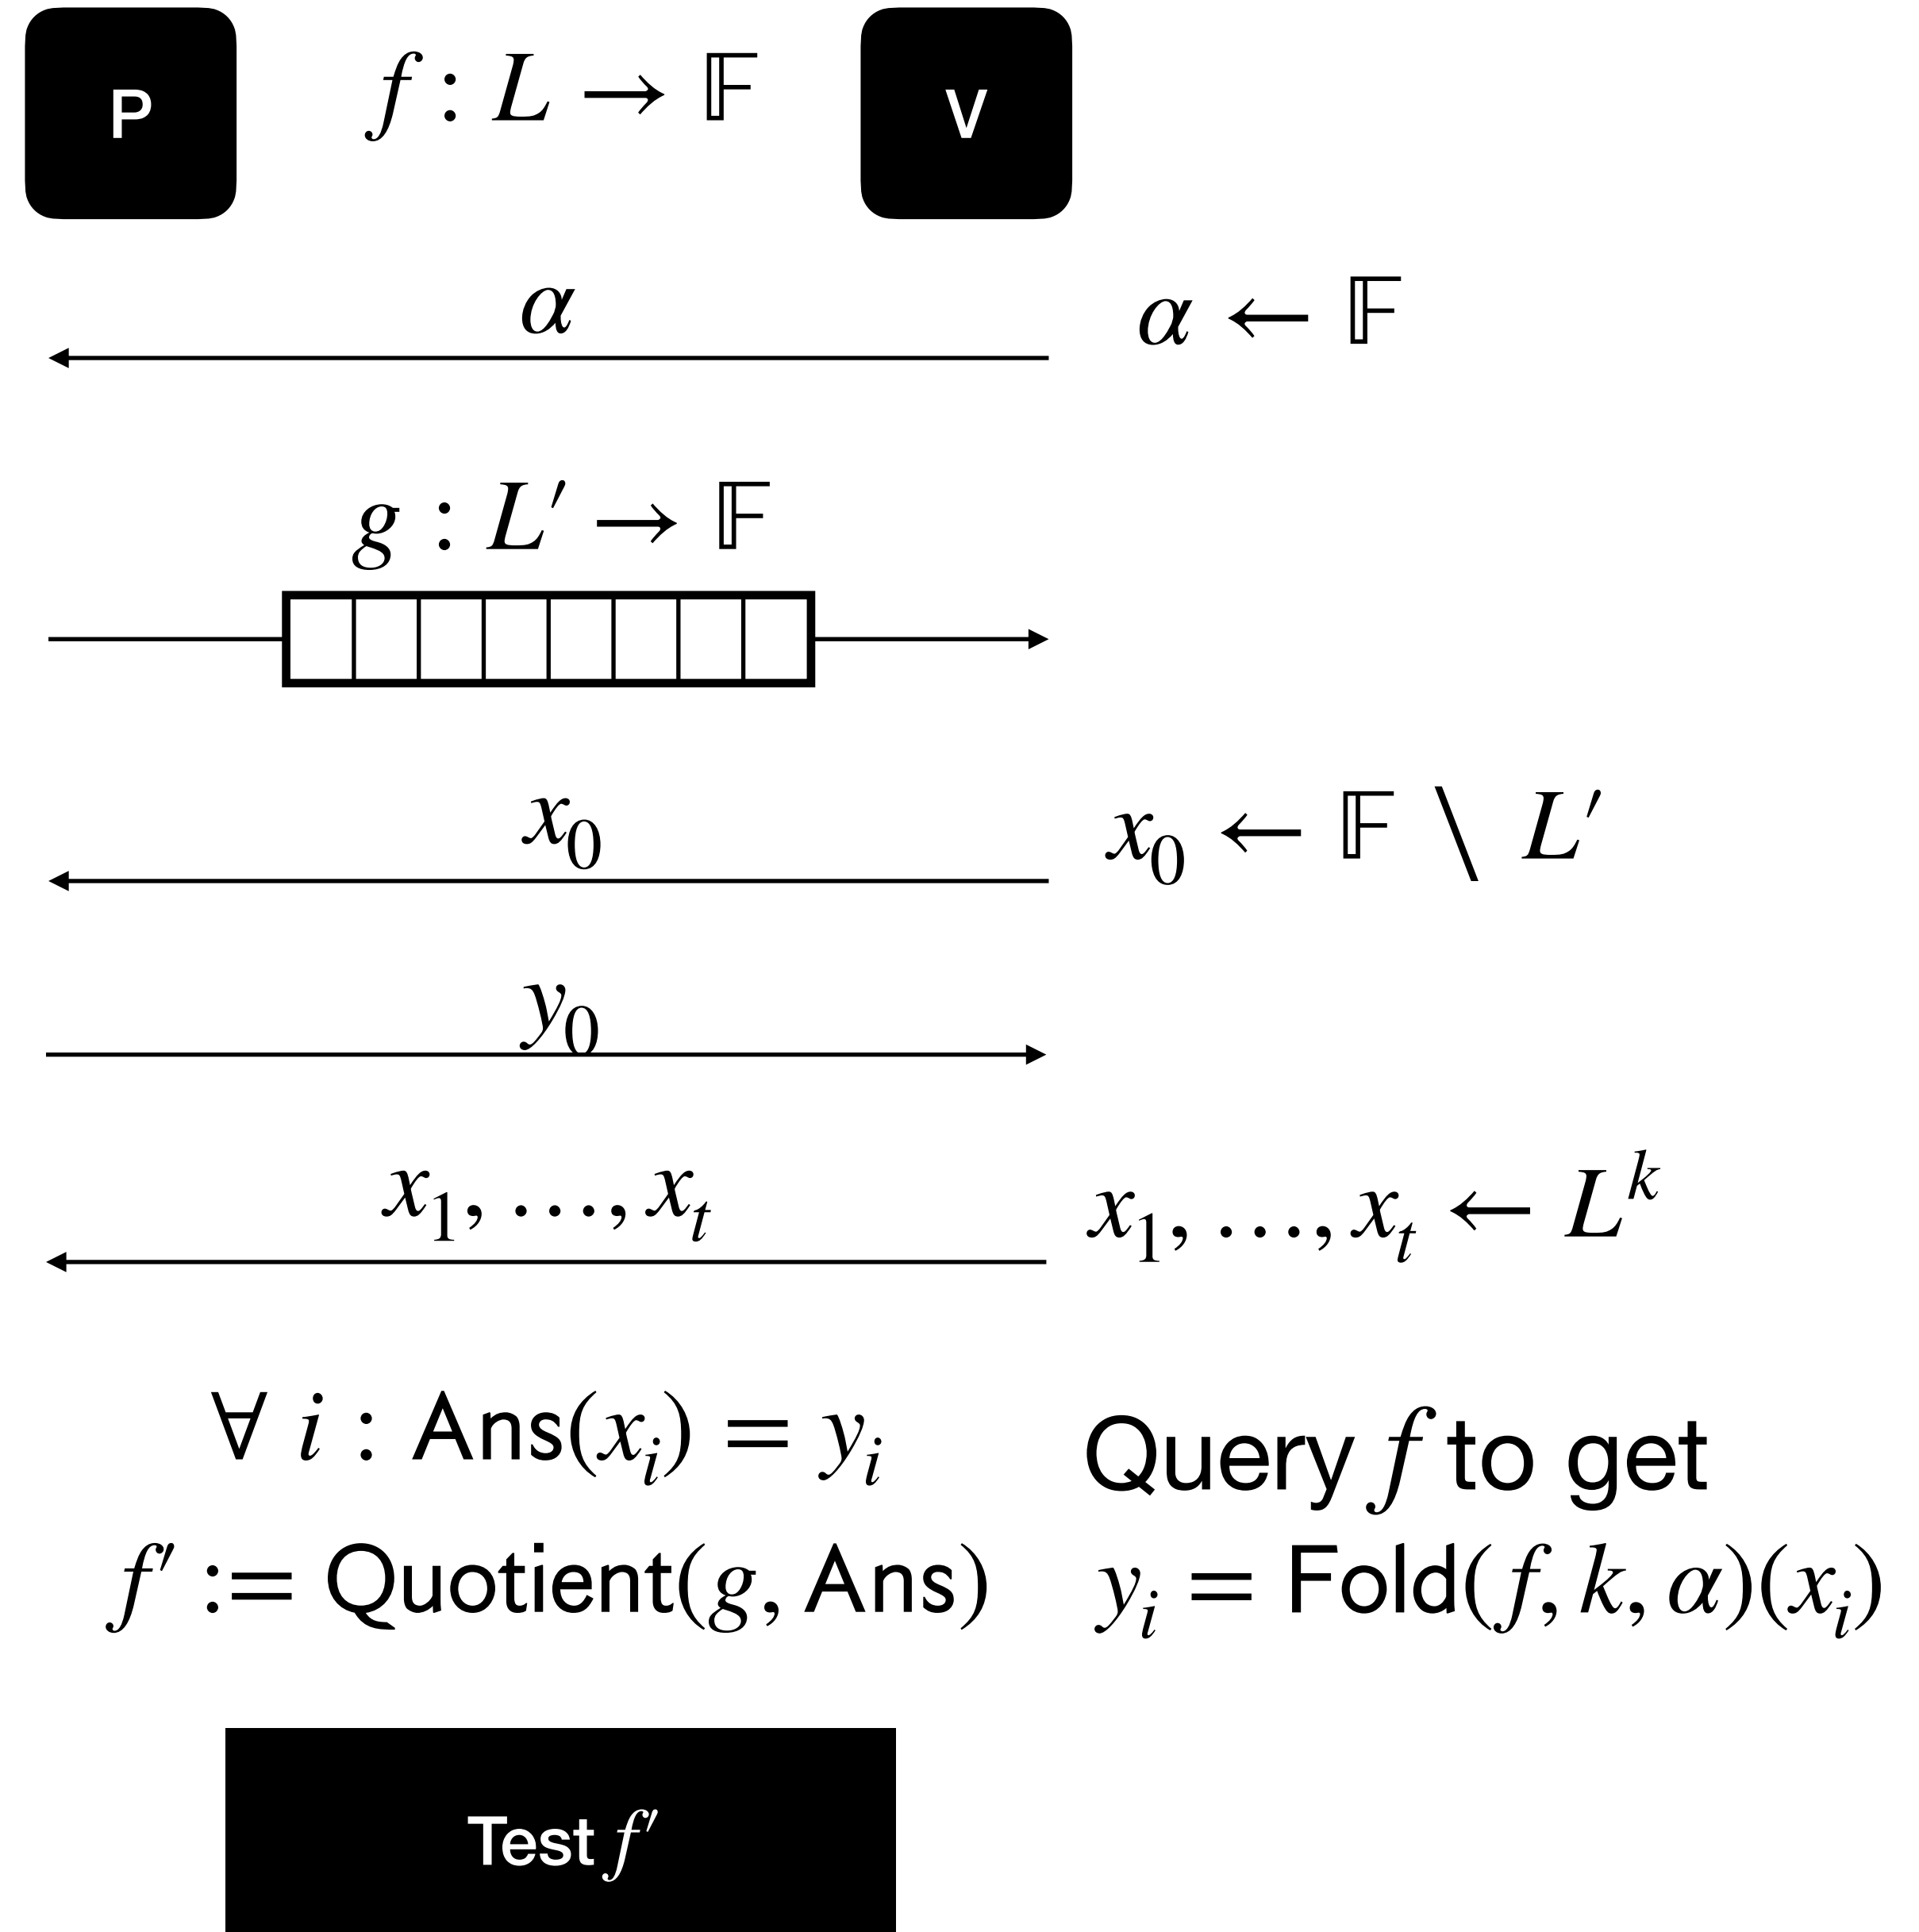
g is claimed to be equal to (the extension of) $\text{Fold}(f, k, \alpha)$ on L'

Problem: We can only query $\text{Fold}(f, k, \alpha)$ on $L^k \neq L'$.

Enforce consistency via Quotient!

Query $\text{Fold}(f, k, \alpha)$ at $x_1, \dots, x_t \in L^k$ to get y_1, \dots, y_t

New function is quotient of g w.r.t. to these points + OOD sample

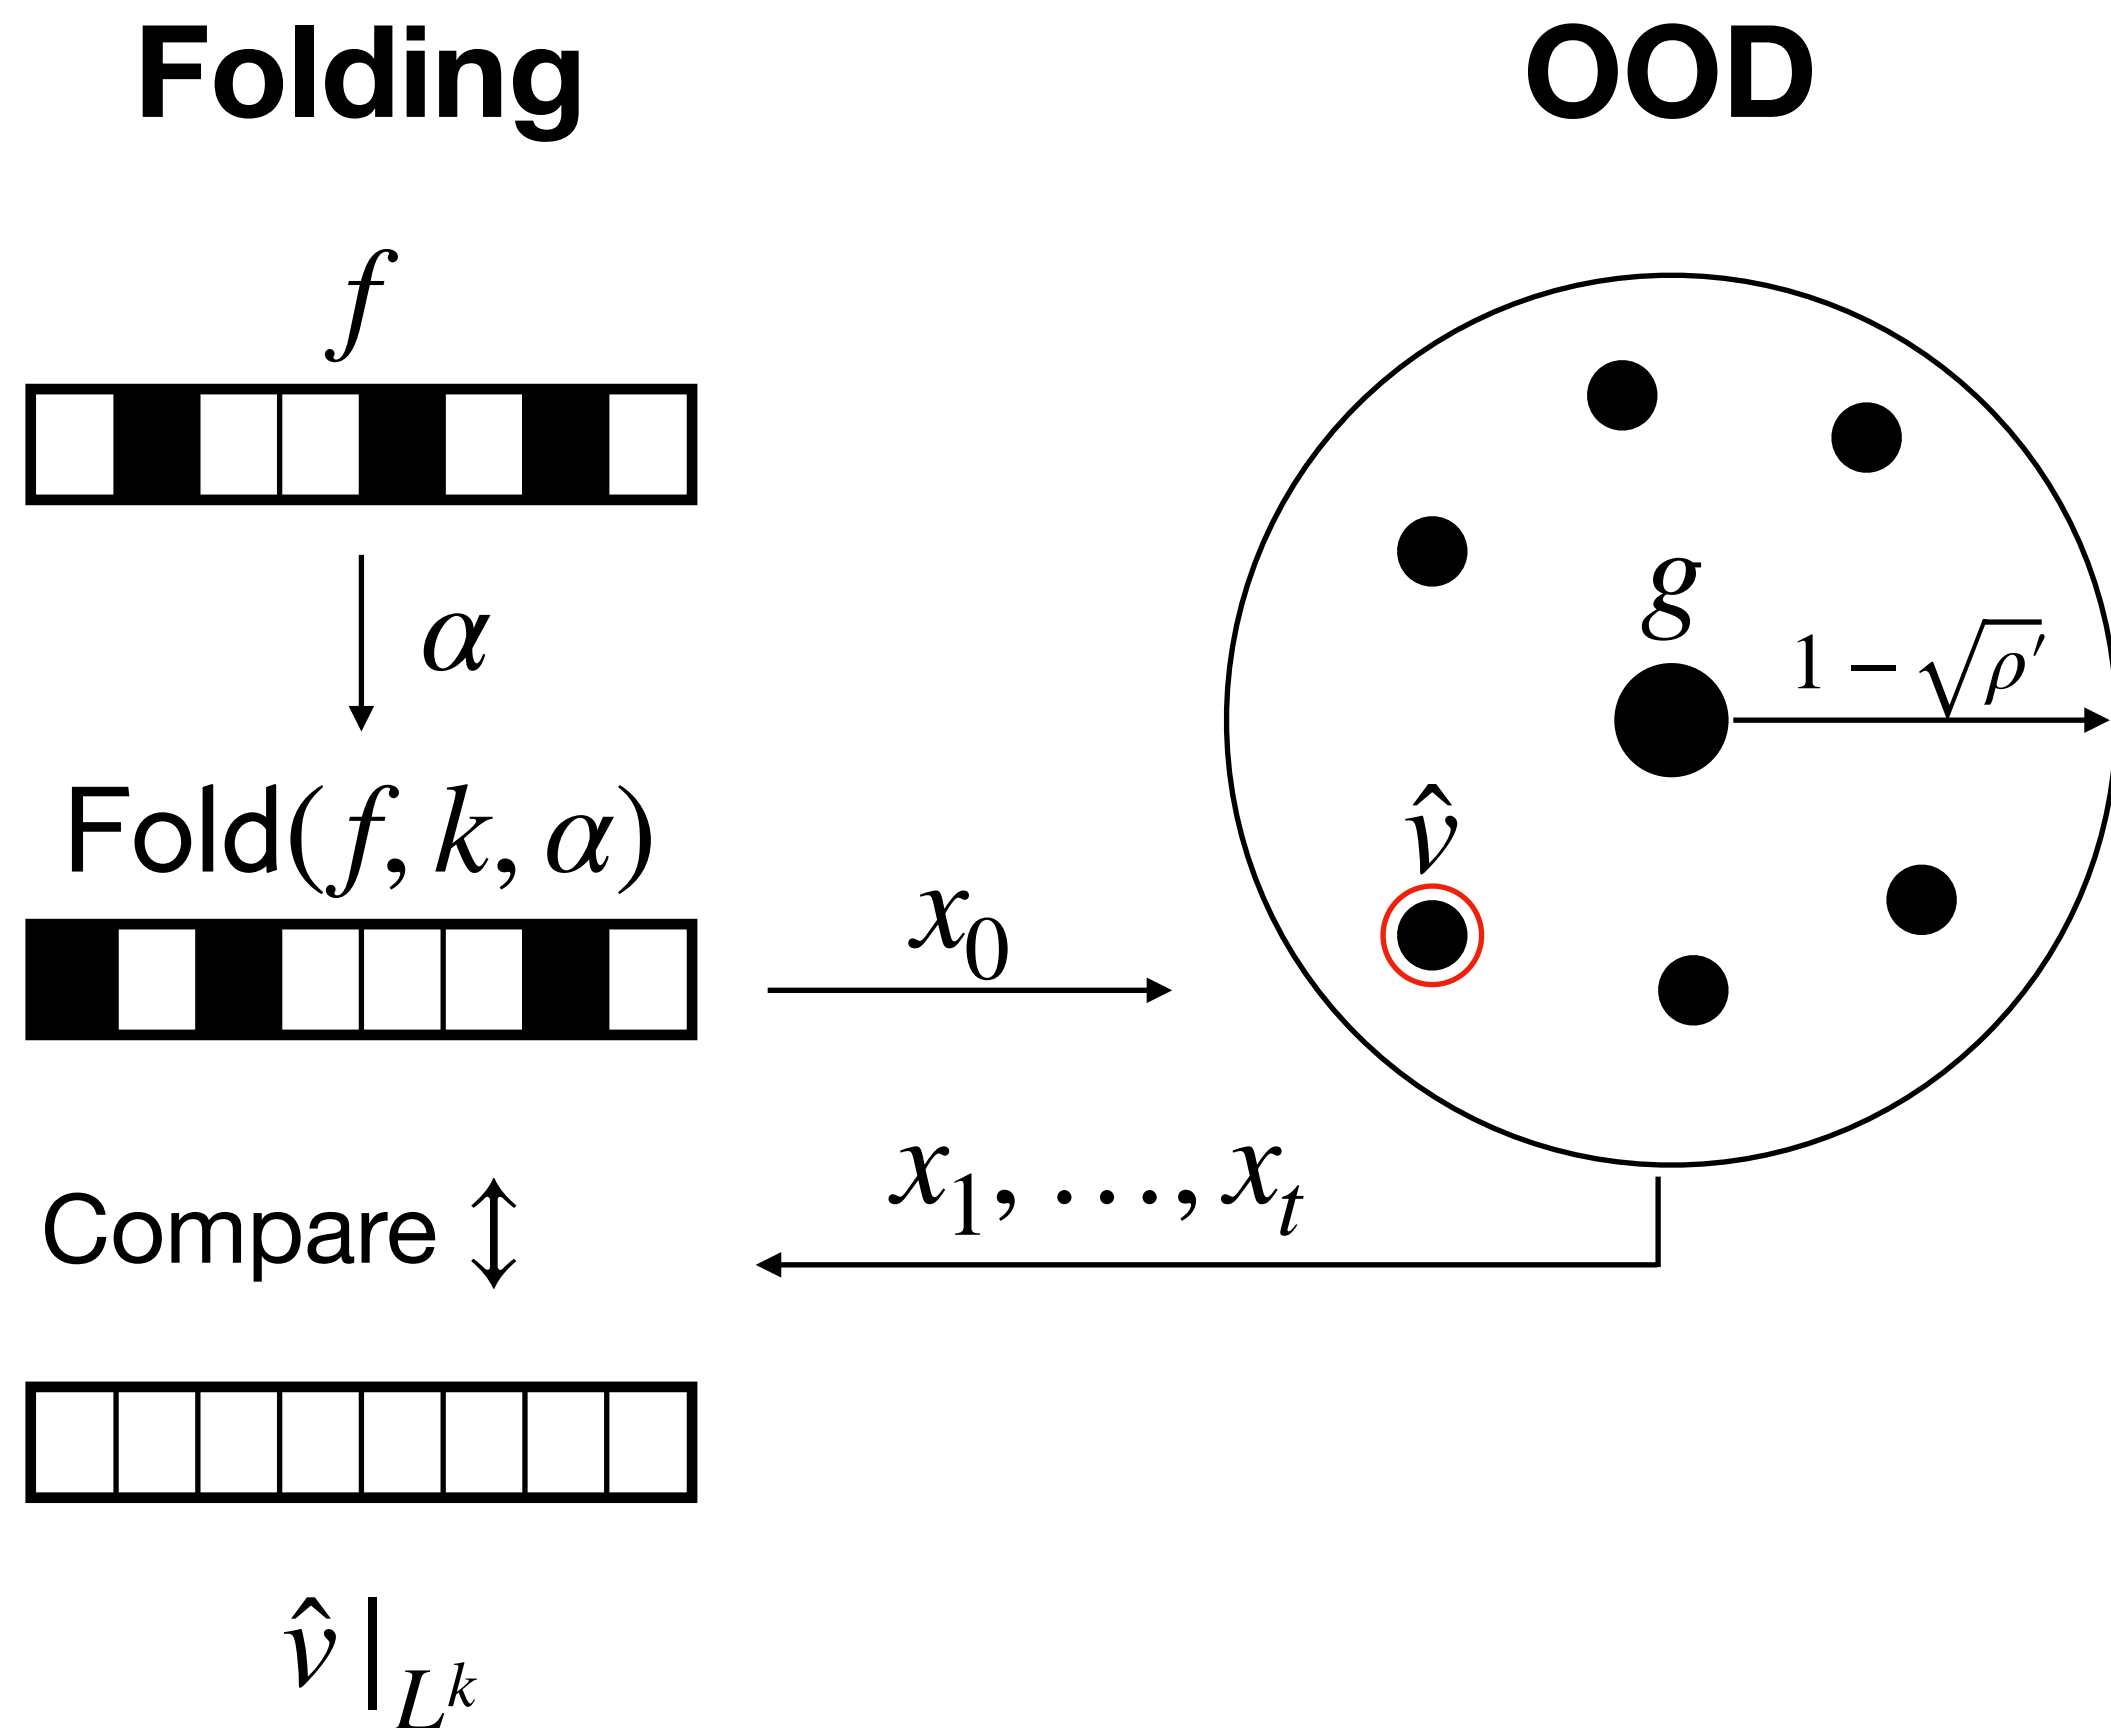


Soundness Analysis

Claim: if f is δ -far from C , unless with probability $\approx (1 - \delta)^t$, f' is $(1 - \sqrt{\rho'})$ far from C'

\hat{v} is **unique** close codeword to g with $\hat{v}(x_0) = y_0$

If at any point $\hat{v}(x_i) \neq y_i$ then, by **quotients**, f' is $(1 - \sqrt{\rho'})$ -far from C'




Since Fold is δ -far from the code,
 $\Delta(\hat{v} \upharpoonright_{L^k}, \text{Fold}(f, k, \alpha)) > \delta$

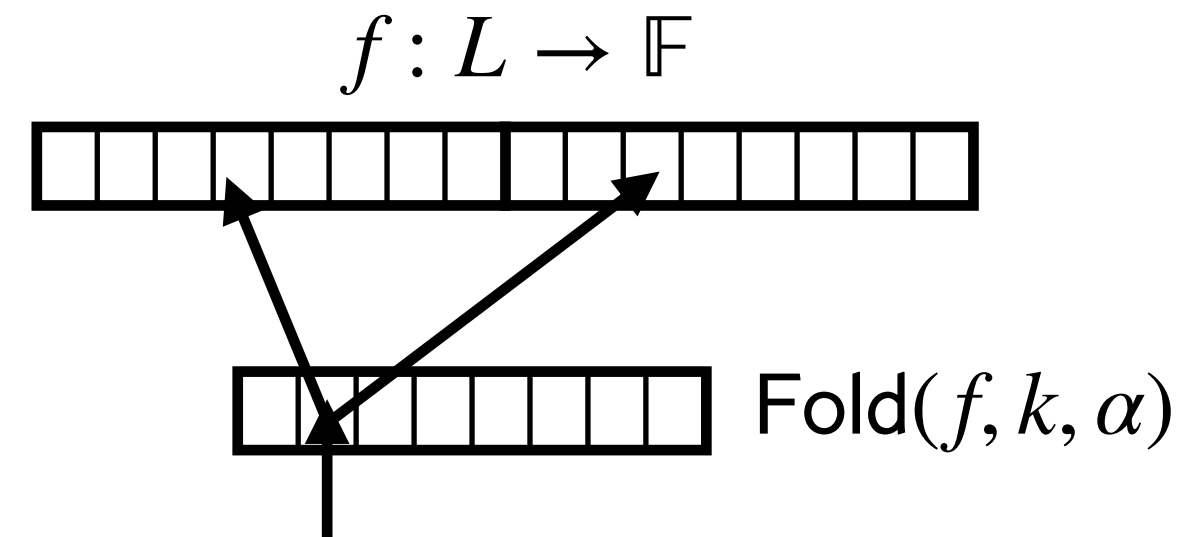
$$\begin{aligned} & \Pr \left[f' \text{ is } 1 - \sqrt{\rho'} \text{ close} \right] \\ & \leq \Pr \left[\forall i, \hat{v}(x_i) = y_i \right] \\ & = \Pr \left[\forall i, \hat{v}(x_i) = \text{Fold}(f, k, \alpha)(x_i) \right] \\ & \leq (1 - \delta)^t \end{aligned}$$

Conclusion

What we saw

What we did have time to talk about

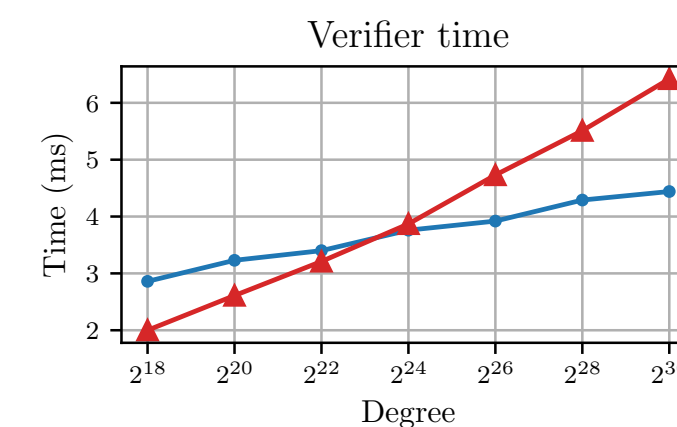
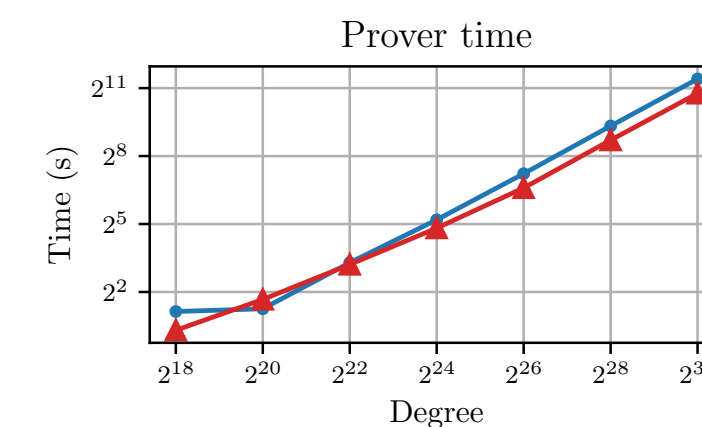
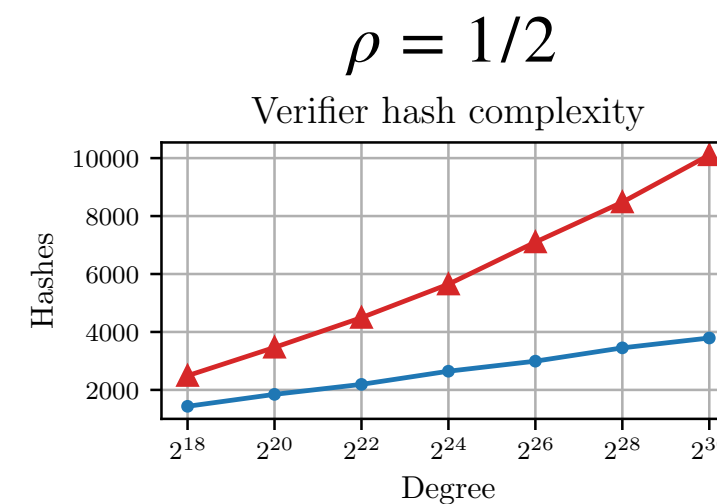



- Techniques
 - Folding and its properties
 - Quotienting and its properties
 - Out-Of-Domain sampling
- STIR 
- Soundness analysis of STIR



```
pub trait LowDegreeTest<F, MerkleConfig, FSConfig>
  where
    F: FftField,
    MerkleConfig: Config,
    FSConfig: CryptographicSponge,
    FSConfig::Config: Clone,
  {
    type Prover: Prover<
      F,
      MerkleConfig,
      FSConfig,
      Commitment = <Self::Verifier as Verifier<F, MerkleConfig, FSConfig>>::Commitment,
      Proof = <Self::Verifier as Verifier<F, MerkleConfig, FSConfig>>::Proof,
    >;
    type Verifier: Verifier<F, MerkleConfig, FSConfig>;

    fn instantiate(
      parameters: Parameters<F, MerkleConfig, FSConfig>,
    ) -> (Self::Prover, Self::Verifier) {
      let prover = Self::Prover::new(parameters.clone());
      let verifier = Self::Verifier::new(parameters);
      (prover, verifier)
    }
  }

```



There is more!

What we did not have time to talk about

Degree corrections

- $\text{Quotient}(f, \text{Ans})$ has degree $d - |S|$, how to bump up to d ?

High-soundness compiler for Poly-IOPs

- Builds on compiler in [ACY23] to achieve concrete efficiency

Round-by-round soundness of STIR \implies secure in non-interactive setting

What's next?

What's next? time to talk about next talk!

• On David, Shahar 🙄)

• Basefold Sm


• **Breaking the $O(\log d)$ -q**

• Work in [ACY23] achieves $O(\log \log$

• Not concretely efficient, lacks efficient soundness

• Exciting!!!

USE STIR



Thank you!

See paper:
ia.cr/2024/390



And blog post:
gfenzi.io/papers/stir



Extra slides

What about the conjecture?

FRI and STIR benefit in roughly the same way

- Conjecture on list-decoding up to distance $1 - \rho$ (instead of $1 - \sqrt{\rho}$)
- FRI queries:

$$O\left(\lambda \cdot \frac{\log d}{-\log \rho}\right)$$

In both, for $\delta = 1 - \rho$,
reduces queries by $\sim 2x$

- STIR queries:

$$O\left(\lambda \cdot \log\left(\frac{\log d}{-\log \rho}\right) + \log d\right)$$