Fiat-Shamir Security of FRI and Related SNARKs

Alexander R. Block[1,2] Albert Garreta[3] Jonathan Katz[2]
Justin Thaler[1,4] Pratyush Ranjan Tiwari[5] Michał Zając[3]

[1]Georgetown Univeristy

[2]University of Maryland

[3]Nethermind

[4]a16z crypto research

[5]Johns Hopkins University

SNARKs

Succinct **N**on-interactive **AR**guments of **K**nowledge

Succinct **N**on-interactive **AR**guments of **K**nowledge

Prover P

Verifier V

Succinct **N**on-interactive **AR**guments of **K**nowledge

$\boxed{\mathcal{L} \in \mathbf{NP}}$

Prover P

Verifier V

Succinct **N**on-interactive **AR**guments of **K**nowledge

$\mathcal{L} \in \mathbf{NP}$

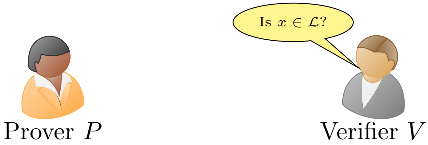

Succinct **N**on-interactive **AR**guments of **K**nowledge

$\mathcal{L} \in \mathbf{NP}$

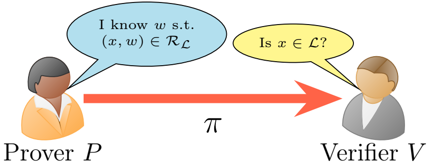

Succinct **N**on-interactive **AR**guments of **K**nowledge

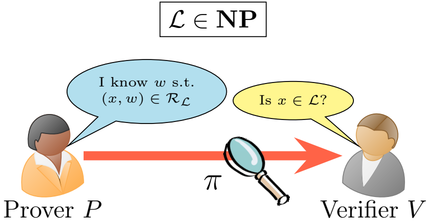

Succinct **N**on-interactive **AR**guments of **K**nowledge

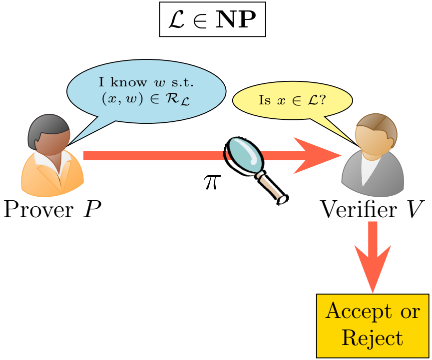

Succinct **N**on-interactive **AR**guments of **K**nowledge

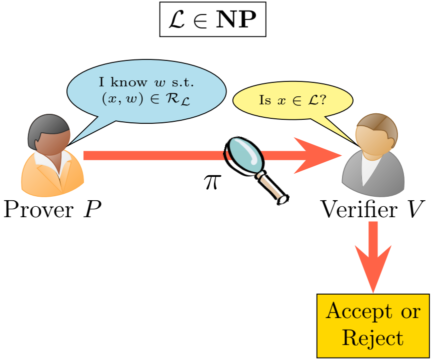

Completeness: $\forall (x, w) \in \mathcal{R}_\mathcal{L}$:

$$\Pr[V(x, \pi) = 1 \mid \pi \leftarrow P(x, w)] = 1$$

Succinct **N**on-interactive **AR**guments of **K**nowledge

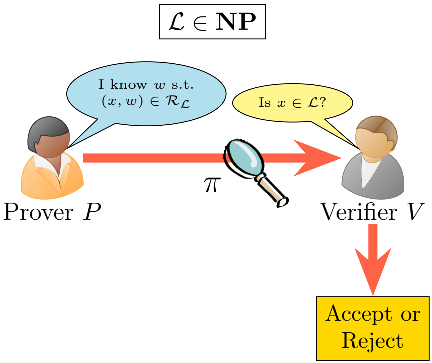

Completeness: $\forall (x, w) \in \mathcal{R}_\mathcal{L}$:

$$\Pr[V(x, \pi) = 1 \mid \pi \leftarrow P(x, w)] = 1$$

ε**-Soundness:** $\forall x \notin \mathcal{L}$, $\forall$ PPT P^*:

$$\Pr[V(x, \pi^*) = 1 \mid \pi^* \stackrel{\$}{\leftarrow} P^*(x)] \leqslant \varepsilon(x, \lambda)$$

Succinct Non-interactive ARguments of Knowledge

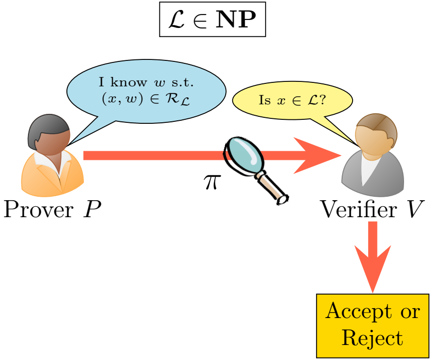

Completeness: $\forall (x, w) \in \mathcal{R}_\mathcal{L}$:

$$\Pr[V(x, \pi) = 1 \mid \pi \leftarrow P(x, w)] = 1$$

ε-Soundness: $\forall x \notin \mathcal{L}$, $\forall$ PPT P^*:

$$\Pr[V(x, \pi^*) = 1 \mid \pi^* \xleftarrow{\$} P^*(x)] \leqslant \varepsilon(x, \lambda)$$

ε-Knowledge Soundness: $\exists$ PPT extractor $\mathcal{E}$ such that $\forall x$ and $\forall$ PPT P^*:

$$\Pr[(x, \mathcal{E}^{P^*}(x)) \in \mathcal{R}_\mathcal{L}] + \varepsilon(x, \lambda) \geqslant \Pr[V(x, \pi^*) = 1 \mid \pi^* \xleftarrow{\$} P^*(x)]$$

Succinct **N**on-interactive **AR**guments of **K**nowledge

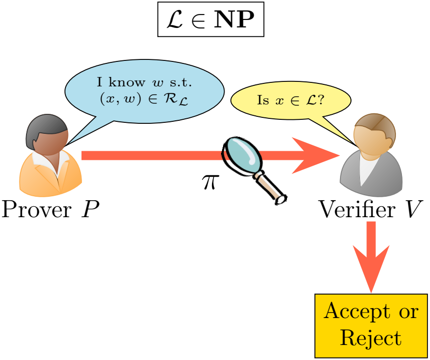

Completeness: $\forall (x, w) \in \mathcal{R}_{\mathcal{L}}$:

$$\Pr[V(x, \pi) = 1 \mid \pi \leftarrow P(x, w)] = 1$$

ε-Soundness: $\forall x \notin \mathcal{L}$, $\forall$ PPT P^*:

$$\Pr[V(x, \pi^*) = 1 \mid \pi^* \xleftarrow{\$} P^*(x)] \leqslant \varepsilon(x, \lambda)$$

ε-Knowledge Soundness: $\exists$ PPT extractor $\mathcal{E}$ such that $\forall x$ and $\forall$ PPT P^*:

$$\Pr[(x, \mathcal{E}^{P^*}(x)) \in \mathcal{R}_{\mathcal{L}}] + \varepsilon(x, \lambda) \geqslant \Pr[V(x, \pi^*) = 1 \mid \pi^* \xleftarrow{\$} P^*(x)]$$

Succinctness: $|\pi| = o_\lambda(|w|)$; ideally $O_\lambda(\text{polylog}(|w|))$

SNARK Construction Paradigm

SNARK Construction Paradigm

1. Construct a public-coin Interactive Oracle Proof (IOP) for $\mathcal{L} \in \mathbf{NP}$

SNARK Construction Paradigm

1. Construct a public-coin Interactive Oracle Proof (IOP) for $\mathcal{L} \in \mathbf{NP}$

$\mathcal{L} \in \mathbf{NP}$

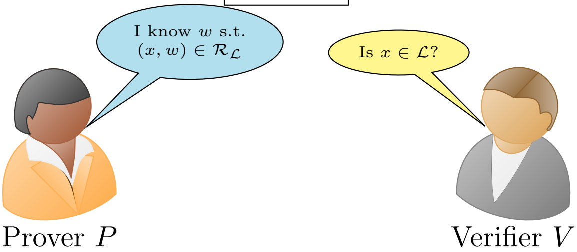

SNARK Construction Paradigm

1. Construct a public-coin Interactive Oracle Proof (IOP) for $\mathcal{L} \in \mathbf{NP}$

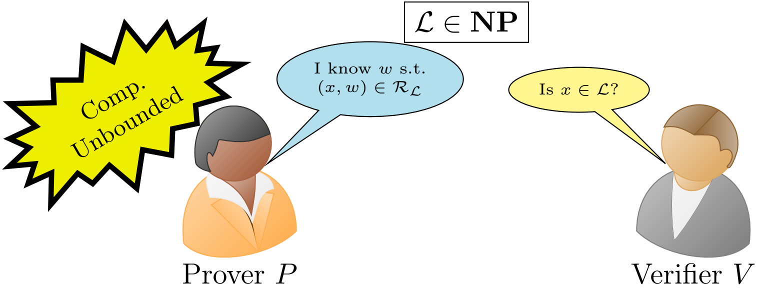

1. Construct a public-coin Interactive Oracle Proof (IOP) for $\mathcal{L} \in \mathbf{NP}$

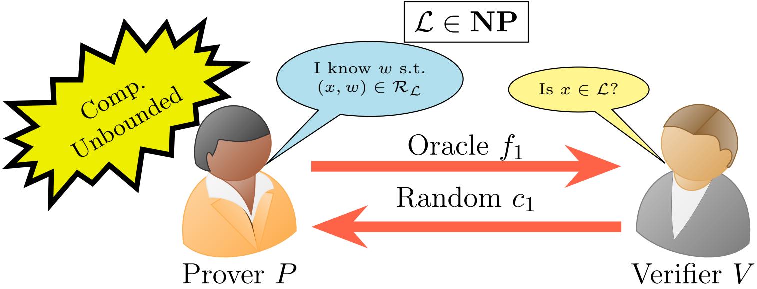

1. Construct a public-coin Interactive Oracle Proof (IOP) for $\mathcal{L} \in \mathbf{NP}$

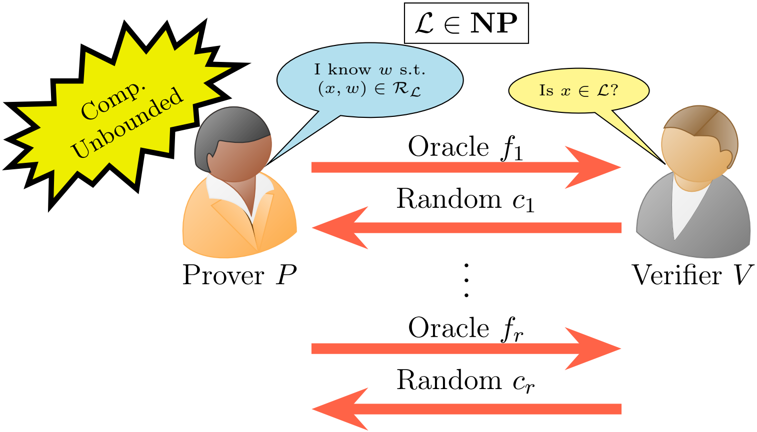

SNARK Construction Paradigm

1. Construct a public-coin Interactive Oracle Proof (IOP) for $\mathcal{L} \in \mathbf{NP}$

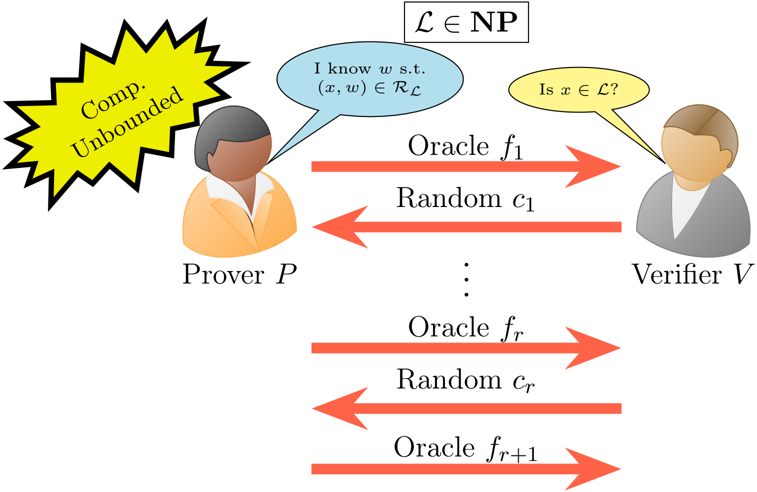

SNARK Construction Paradigm

1. Construct a public-coin Interactive Oracle Proof (IOP) for $\mathcal{L} \in \mathbf{NP}$

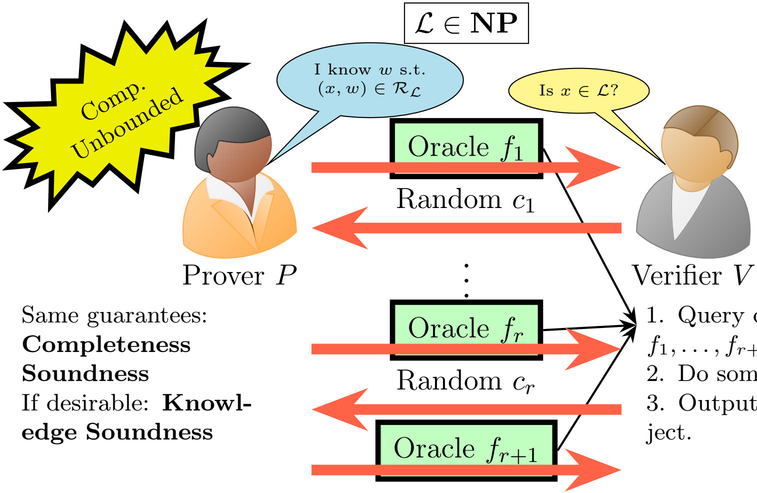

Same guarantees: **Completeness** **Soundness** If desirable: **Knowledge Soundness**

1. Query oracles $f_1, \ldots, f_{r+1}$.
2. Do some computation.
3. Output Accept or Reject.

SNARK Construction Paradigm

2. Replace oracles with Merkle trees, and replace Verifier queries with Merkle authentication paths

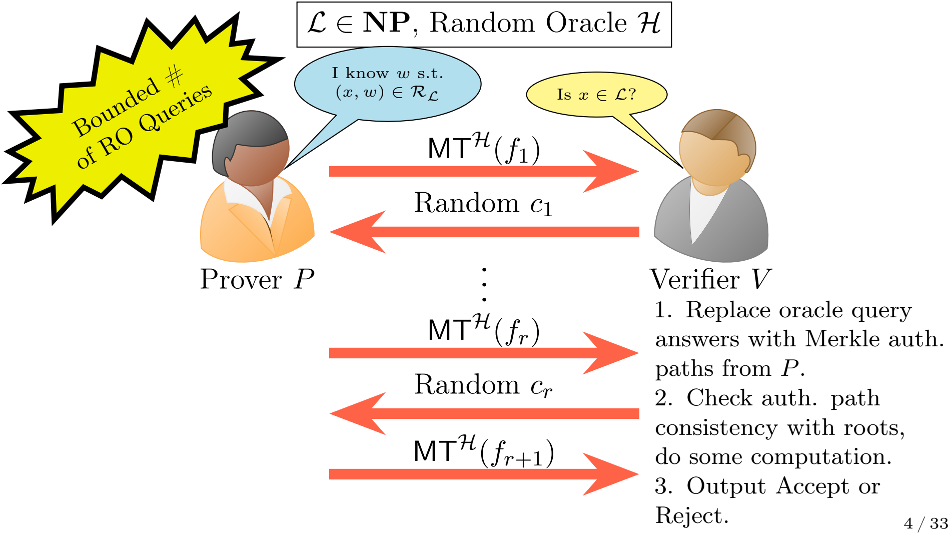

SNARK Construction Paradigm

3. Compress Merkle tree protocol with Fiat-Shamir by replacing V challenges with output of $\mathcal{H}$

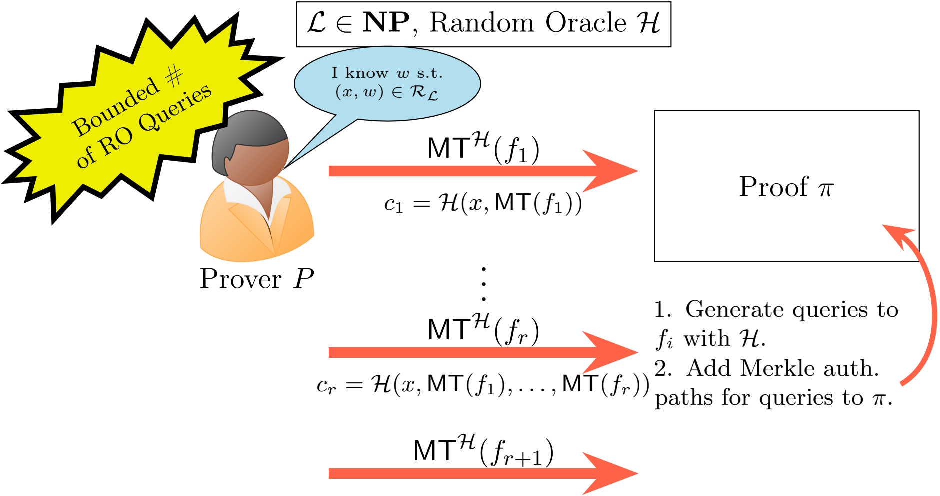

Security of Fiat-Shamir Transformation

- Not secure in general [Bar01, GK03, BDG$^+$13], even in RO model, for many-round ($\omega(1)$-round) protocols
  - E.g., sequential/parallel repetition of constant-sound interactive protocols

Security of Fiat-Shamir Transformation

- Not secure in general [Bar01, GK03, BDG+13], even in RO model, for many-round ($\omega(1)$-round) protocols
  - E.g., sequential/parallel repetition of constant-sound interactive protocols
- FS often applied to many-round protocols **without** formal security proofs
  - Often only prove *interactive security*

Our Results: Bird's Eye View

This Work

Our Results: Bird's Eye View

1. Prove FS security of the **FRI Protocol** [BBH+18] and the **batched FRI Protocol**

This Work

Our Results: Bird's Eye View

1. Prove FS security of the **FRI Protocol** [BBH+18] and the **batched FRI Protocol**
   - Fills security gaps in [CMS19, COS20, KPV22]

This Work

Our Results: Bird's Eye View

1. Prove FS security of the **FRI Protocol** [BBH+18] and the **batched FRI Protocol**
   - Fills security gaps in [CMS19, COS20, KPV22]

2. Introduce δ**-Correlated IOPs** and prove their FS security

Our Results: Bird's Eye View

1. Prove FS security of the **FRI Protocol** [BBH+18] and the **batched FRI Protocol**
   - Fills security gaps in [CMS19, COS20, KPV22]

2. Introduce **δ-Correlated IOPs** and prove their FS security
   - Intuitively, these are protocols that use FRI as a sub-routine

This Work

Our Results: Bird's Eye View

1. Prove FS security of the **FRI Protocol** [BBH+18] and the **batched FRI Protocol**
   - Fills security gaps in [CMS19, COS20, KPV22]
2. Introduce δ**-Correlated IOPs** and prove their FS security
   - Intuitively, these are protocols that use FRI as a sub-routine
3. Formulate a δ-Correlated IOP which captures many "Plonk-like" protocols and prove their FS security

This Work

Our Results: Bird's Eye View

1. Prove FS security of the **FRI Protocol** [BBH+18] and the **batched FRI Protocol**
   - Fills security gaps in [CMS19, COS20, KPV22]
2. Introduce δ**-Correlated IOPs** and prove their FS security
   - Intuitively, these are protocols that use FRI as a sub-routine
3. Formulate a δ-Correlated IOP which captures many "Plonk-like" protocols and prove their FS security
   - Captures Plonky2 [Pol], Redshift [KPV22], RISC Zero [Tea23]
   - ethSTARK [Sta23] and DEEP-FRI [BGK+20] also fit in this framework

This Work

Our Results: Bird's Eye View

1. Prove FS security of the **FRI Protocol** [BBH+18] and the **batched FRI Protocol**
   - Fills security gaps in [CMS19, COS20, KPV22]
2. Introduce δ-**Correlated IOPs** and prove their FS security
   - Intuitively, these are protocols that use FRI as a sub-routine
3. Formulate a δ-Correlated IOP which captures many "Plonk-like" protocols and prove their FS security
   - Captures Plonky2 [Pol], Redshift [KPV22], RISC Zero [Tea23]
   - ethSTARK [Sta23] and DEEP-FRI [BGK+20] also fit in this framework

"Plonk-like" $\approx$ protocols which use FRI + a permutation argument [Lip89, Lip90, ZGK+18, BEG+94, BCG+18], helped popularized by the PLONK SNARK [GWC19]

Context: Why FRI?

Context: Why FRI?

- FRI is popular in the SNARK space

Context: Why FRI?

- FRI is popular in the SNARK space

80-bits of **conjectured** interactive security

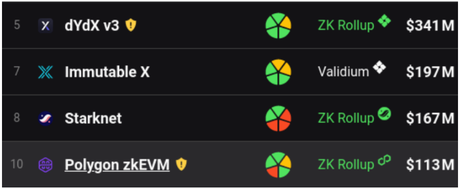

5	dYdX v3	ZK Rollup	$341 M
7	Immutable X	Validium	$197 M
8	Starknet	ZK Rollup	$167 M
10	Polygon zkEVM	ZK Rollup	$113 M

96-bits of **conjectured** interactive security

Also uses FRI

Context: Why FRI?

- FRI is popular in the SNARK space

80-bits of **conjectured** interactive security

Also uses FRI

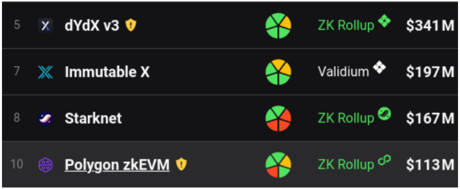

5	dYdX v3	ZK Rollup	$341 M
7	Immutable X	Validium	$197 M
8	Starknet	ZK Rollup	$167 M
10	Polygon zkEVM	ZK Rollup	$113 M

96-bits of **conjectured** interactive security

- Plonk-like protocols are also used in many L2 projects; e.g., [Min, Mat, Suc, Dus, nil]

Before this work, no formal FS security analysis of FRI existed

The FRI Protocol

FRI = **F**ast **R**eed-Solomon **I**OP of Proximity [BBH+18]

The FRI Protocol

FRI = **F**ast **R**eed-Solomon **I**OP of Proximity [BBH+18]

Parameters:

- Finite field $\mathbb{F}$ and $L_0 \subset \mathbb{F}$ of size 2^n
  - L_0 is a smooth multiplicative subgroup

The FRI Protocol

FRI = **F**ast **R**eed-Solomon **I**OP of Proximity [BBH+18]

Parameters:

- Finite field $\mathbb{F}$ and $L_0 \subset \mathbb{F}$ of size 2^n
  - L_0 is a smooth multiplicative subgroup
- Degree bound $d_0 = 2^k$

The FRI Protocol

FRI = **F**ast **R**eed-Solomon **I**OP of Proximity [BBH+18]

Parameters:

- Finite field $\mathbb{F}$ and $L_0 \subset \mathbb{F}$ of size 2^n
  - L_0 is a smooth multiplicative subgroup
- Degree bound $d_0 = 2^k$
- $\mathsf{RS}^0 := \mathsf{RS}[\mathbb{F}, L_0, d_0] = \{(f(z))_{z \in L_0} \colon f(X) \in \mathbb{F}^{<d_0}[X]\}$

The FRI Protocol

FRI = **F**ast **R**eed-Solomon **I**OP of Proximity [BBH+18]

Parameters:

- Finite field $\mathbb{F}$ and $L_0 \subset \mathbb{F}$ of size 2^n
    - L_0 is a smooth multiplicative subgroup
- Degree bound $d_0 = 2^k$
- $\mathsf{RS}^0 := \mathsf{RS}[\mathbb{F}, L_0, d_0] = \{(f(z))_{z \in L_0} : f(X) \in \mathbb{F}^{<d_0}[X]\}$
- Rate $\rho = d_0/|L_0| = 2^{-(n-k)}$, proximity parameter $\delta \in (0, 1 - \sqrt{\rho})$, verifier repetition parameter $\ell \in \mathbb{Z}^+$

The FRI Protocol

FRI = **F**ast **R**eed-Solomon **I**OP of Proximity [BBH+18]

Parameters:

- Finite field $\mathbb{F}$ and $L_0 \subset \mathbb{F}$ of size 2^n
  - L_0 is a smooth multiplicative subgroup
- Degree bound $d_0 = 2^k$
- $\mathsf{RS}^0 := \mathsf{RS}[\mathbb{F}, L_0, d_0] = \{(f(z))_{z \in L_0} \colon f(X) \in \mathbb{F}^{<d_0}[X]\}$
- Rate $\rho = d_0/|L_0| = 2^{-(n-k)}$, proximity parameter $\delta \in (0, 1 - \sqrt{\rho})$, verifier repetition parameter $\ell \in \mathbb{Z}^+$

FRI proves that a function $G_0 \colon L_0 \to \mathbb{F}$ is δ-close to RS^0

- **Round-by-round (Knowledge) Soundness** [CCH+19, CMS19]

RBR Soundness: Intuition

Showing Fiat-Shamir Security

- **Round-by-round (Knowledge) Soundness** [CCH+19, CMS19]

RBR Soundness: Intuition

If $x \notin \mathcal{L}$, then protocol is "**doomed**"

No matter what the prover does, the protocol should forever remain "**doomed**"

Round-by-round Soundness

Definition 1 (RBR Soundness)

A protocol Π for a language $\mathcal{L}$ has RBR soundness error ε if $\exists$ a "doomed" set of (partial) transcripts $\mathcal{D}$ such that:

Definition 1 (RBR Soundness)

A protocol Π for a language $\mathcal{L}$ has RBR soundness error ε if $\exists$ a "doomed" set of (partial) transcripts $\mathcal{D}$ such that:

1. If $x \notin \mathcal{L}$ then $(x, \emptyset) \in \mathcal{D}$;

Round-by-round Soundness

Definition 1 (RBR Soundness)

A protocol Π for a language $\mathcal{L}$ has RBR soundness error ε if $\exists$ a "doomed" set of (partial) transcripts $\mathcal{D}$ such that:

1. If $x \notin \mathcal{L}$ then $(x, \emptyset) \in \mathcal{D}$;
2. For all complete transcripts τ, if $(x, \tau) \in \mathcal{D}$ then the verifier rejects; and

Round-by-round Soundness

Definition 1 (RBR Soundness)

A protocol Π for a language $\mathcal{L}$ has RBR soundness error ε if $\exists$ a "doomed" set of (partial) transcripts $\mathcal{D}$ such that:

1. If $x \notin \mathcal{L}$ then $(x, \emptyset) \in \mathcal{D}$;
2. For all complete transcripts τ, if $(x, \tau) \in \mathcal{D}$ then the verifier rejects; and
3. If τ_{i-1} is an $(i-1)$-partial transcript and $(x, \tau_{i-1}) \in \mathcal{D}$, then for all prover messages m:

$$\Pr_c[(x, \tau_{i-1} \| m \| c) \notin \mathcal{D}] \leqslant \varepsilon.$$

RBR Soundness and Fiat-Shamir Security

Round-by-round soundness implies Fiat-Shamir security

RBR Soundness and Fiat-Shamir Security

Round-by-round soundness implies Fiat-Shamir security

RO Model
Q-query adversary
κ-bit RO output

RBR Soundness and Fiat-Shamir Security

Round-by-round soundness implies Fiat-Shamir security

RBR Error
ε

RO Model
Q-query adversary
κ-bit RO output

RBR Soundness and Fiat-Shamir Security

Round-by-round soundness implies Fiat-Shamir security

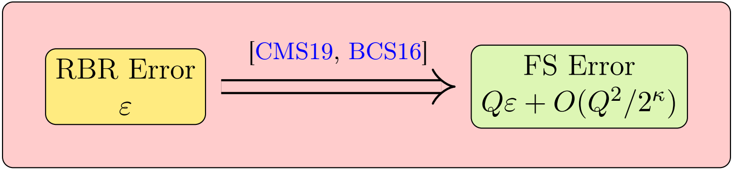

RO Model

Q-query adversary

κ-bit RO output

Our Results: FS Security of FRI

Theorem 1

Let $\mathbb{F}$ be a finite field, $L_0 \subset \mathbb{F}^*$ be a smooth multiplicative subgroup of size 2^n, $d_0 = 2^k$, $\rho = d_0/|L_0|$, and $\ell \in \mathbb{Z}^+$. For any integer $m \geqslant 3$, $\eta \in (0, \sqrt{\rho}/(2m))$, $\delta \in (0, 1 - \sqrt{\rho} - \eta)$, and function $G_0 \colon L_0 \to \mathbb{F}$ that is δ-far from $\mathsf{RS}[\mathbb{F}, L_0, d_0]$, the FRI protocol has RBR (knowledge) soundness error

$$\varepsilon_{\mathsf{rbr}}^{\mathsf{FRI}} = \max\left\{\frac{(m+1/2)^7|L_0|^2}{3\rho^{3/2}|\mathbb{F}|}, (1-\delta)^\ell\right\}.$$

Our Results: FS Security of FRI

Theorem 1

Let $\mathbb{F}$ be a finite field, $L_0 \subset \mathbb{F}^*$ be a smooth multiplicative subgroup of size 2^n, $d_0 = 2^k$, $\rho = d_0/|L_0|$, and $\ell \in \mathbb{Z}^+$. For any integer $m \geqslant 3$, $\eta \in (0, \sqrt{\rho}/(2m))$, $\delta \in (0, 1 - \sqrt{\rho} - \eta)$, and function $G_0 \colon L_0 \to \mathbb{F}$ that is δ-far from $\mathsf{RS}[\mathbb{F}, L_0, d_0]$, the FRI protocol has RBR (knowledge) soundness error

$$\varepsilon_{\mathsf{rbr}}^{\mathsf{FRI}} = \max\left\{ \frac{(m+1/2)^7 |L_0|^2}{3\rho^{3/2}|\mathbb{F}|}, (1-\delta)^\ell \right\}.$$

- Same result holds for **batched FRI**[1]

[1] When batching with distinct challenges; see paper for details.

Our Results: FS Security of FRI

Theorem 1

Let $\mathbb{F}$ be a finite field, $L_0 \subset \mathbb{F}^*$ be a smooth multiplicative subgroup of size 2^n, $d_0 = 2^k$, $\rho = d_0/|L_0|$, and $\ell \in \mathbb{Z}^+$. For any integer $m \geqslant 3$, $\eta \in (0, \sqrt{\rho}/(2m))$, $\delta \in (0, 1 - \sqrt{\rho} - \eta)$, and function $G_0 \colon L_0 \to \mathbb{F}$ that is δ-far from $\mathsf{RS}[\mathbb{F}, L_0, d_0]$, the FRI protocol has RBR (knowledge) soundness error

$$\varepsilon_{\mathsf{rbr}}^{\mathsf{FRI}} = \max\left\{ \frac{(m+1/2)^7 |L_0|^2}{3\rho^{3/2}|\mathbb{F}|}, (1-\delta)^\ell \right\}.$$

- Same result holds for **batched FRI**[1]
- Implies FS error $Q \cdot \varepsilon_{\mathsf{rbr}}^{\mathsf{FRI}} + O(Q^2/2^\kappa)$ in the ROM

[1] When batching with distinct challenges; see paper for details.

Comparison with FRI Interactive Security

- Best provable interactive soundness of FRI [BCI+20] is

$$\varepsilon_1 + \varepsilon_2 + \varepsilon_3, \text{ where}$$

$$\varepsilon_1 = \frac{(m + 1/2)^7 2^{2n}}{3\rho^{3/2}|\mathbb{F}|} \qquad \varepsilon_2 = O\left(\frac{2^n \cdot n}{\sqrt{\rho}|\mathbb{F}|}\right) \qquad \varepsilon_3 = (1-\delta)^\ell$$

Comparison with FRI Interactive Security

- Best provable interactive soundness of FRI [BCI+20] is

$$\varepsilon_1 + \varepsilon_2 + \varepsilon_3, \text{ where}$$

$$\varepsilon_1 = \frac{(m+1/2)^7 2^{2n}}{3\rho^{3/2}|\mathbb{F}|} \qquad \varepsilon_2 = O\left(\frac{2^n \cdot n}{\sqrt{\rho}|\mathbb{F}|}\right) \qquad \varepsilon_3 = (1-\delta)^\ell$$

- We prove FRI has RBR soundness error $= \max\{\varepsilon_1, \varepsilon_3\}$

δ-Correlated IOPs

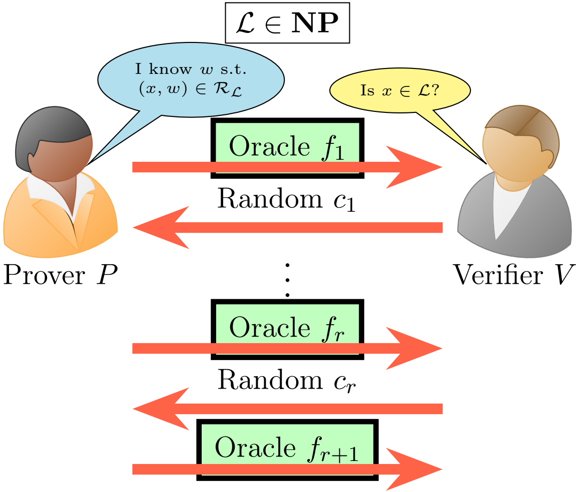

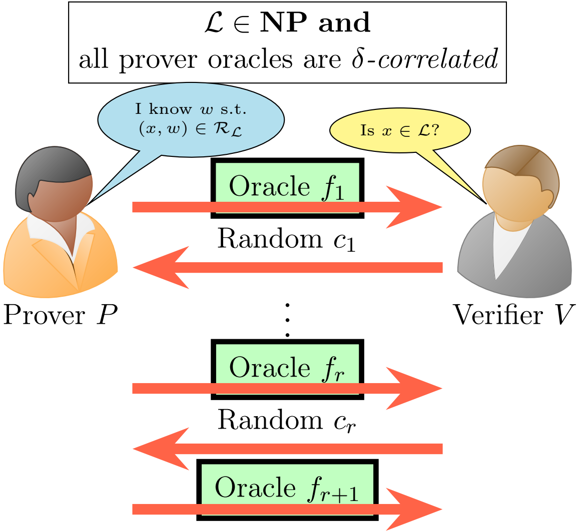
$\mathcal{L} \in$ **NP and**
all prover oracles are δ-*correlated*
I know w s.t. $(x, w) \in \mathcal{R}_{\mathcal{L}}$
Is $x \in \mathcal{L}$?
Oracle f_1
Random c_1
Prover P
Verifier V
Oracle f_r
Random c_r
Oracle f_{r+1}

δ-Correlated IOPs

$\mathcal{L} \in$ NP and
all prover oracles are δ-*correlated*

δ-correlated: for a fixed RS codespace, all f_i are δ-close to RS

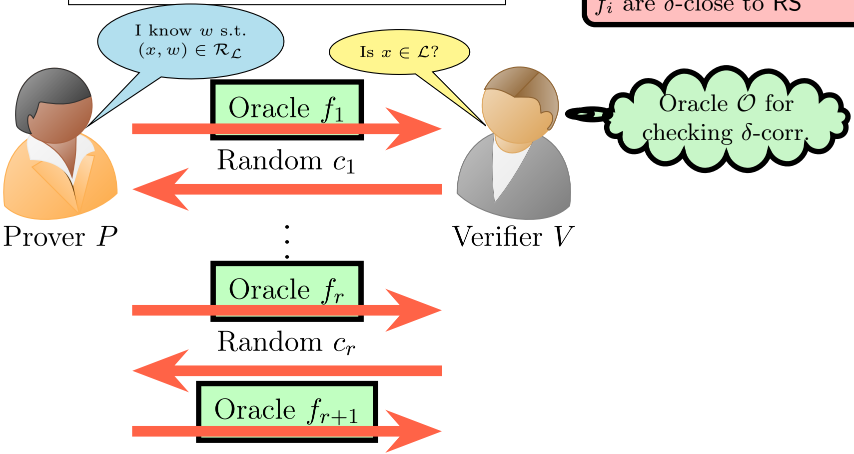

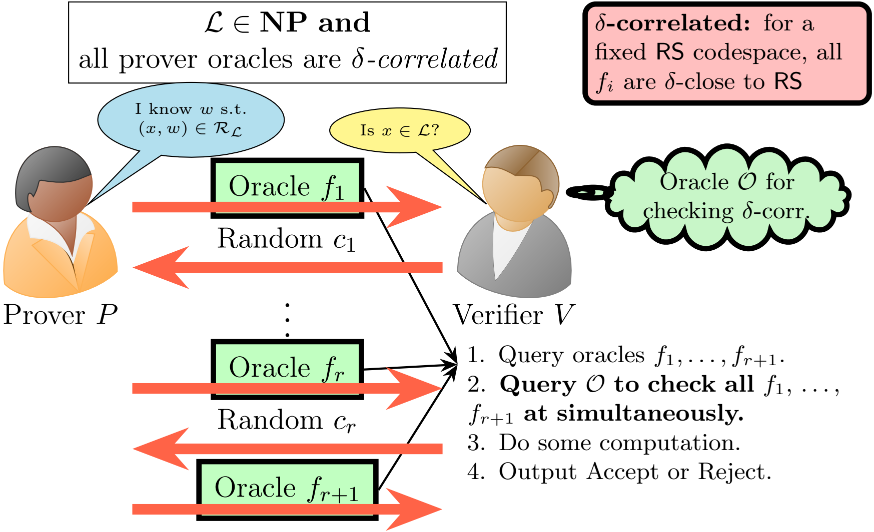
$\mathcal{L} \in$ **NP and**
all prover oracles are δ-*correlated*
δ-correlated: for a fixed RS codespace, all f_i are δ-close to RS
I know w s.t. $(x, w) \in \mathcal{R}_{\mathcal{L}}$
Is $x \in \mathcal{L}$?
Oracle f_1
Random c_1
Oracle $\mathcal{O}$ for checking δ-corr.
Prover P
Verifier V
Oracle f_r
Random c_r
Oracle f_{r+1}
1. Query oracles $f_1, \ldots, f_{r+1}$.
2. **Query $\mathcal{O}$ to check all $f_1, \ldots, f_{r+1}$ at simultaneously.**
3. Do some computation.
4. Output Accept or Reject.

Our Results: δ-Correlated IOPs

Theorem 2 (Informal)

Let $\Pi_\delta^{\mathcal{O}}$ be a δ-correlated IOP for a fixed RS *code of rate $\rho \in (0, 1]$, and let $\eta \in (0, \sqrt{\rho})$.*

Our Results: δ-Correlated IOPs

Theorem 2 (Informal)

Let $\Pi_\delta^{\mathcal{O}}$ be a δ-correlated IOP for a fixed RS *code of rate $\rho \in (0, 1]$, and let $\eta \in (0, \sqrt{\rho})$.*

1. *If $\Pi_0^{\mathcal{O}}$ has RBR (knowledge) error ε, then $\Pi_\delta^{\mathcal{O}}$ has RBR (knowledge) error $\varepsilon/(2\eta\sqrt{\rho})$, where $\delta = 1 - \sqrt{\rho} - \eta > 0$.*

Our Results: δ-Correlated IOPs

Theorem 2 (Informal)

Let $\Pi_\delta^{\mathcal{O}}$ be a δ-correlated IOP for a fixed RS *code of rate $\rho \in (0,1]$, and let $\eta \in (0, \sqrt{\rho})$.*

1. *If $\Pi_0^{\mathcal{O}}$ has RBR (knowledge) error ε, then $\Pi_\delta^{\mathcal{O}}$ has RBR (knowledge) error $\varepsilon/(2\eta\sqrt{\rho})$, where $\delta = 1 - \sqrt{\rho} - \eta > 0$.*
2. *If Π' is an IOP for testing δ-correlation in* RS *with RBR error ε', then $\Pi_\delta^{\Pi'}$ is an IOP with RBR (knowledge) error $\max\{\varepsilon/(2\eta\sqrt{\rho}), \varepsilon'\}$.*

Theorem 2 gives a new paradigm for SNARK design

Theorem 2 gives a new paradigm for SNARK design

RBR (knowledge) sound 0-correlated IOP

Theorem 2 gives a new paradigm for SNARK design

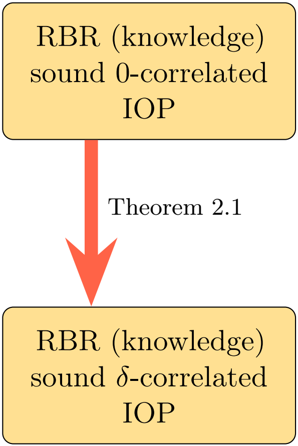

Theorem 2 gives a new paradigm for SNARK design

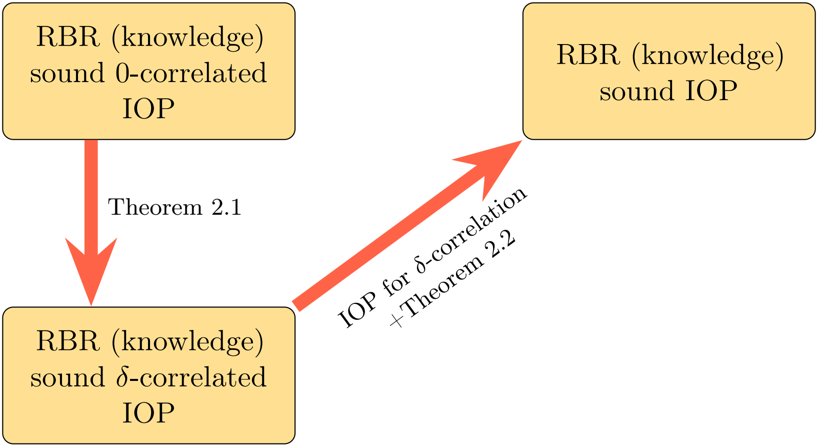

SNARKs from δ-Correlated IOPs

Theorem 2 gives a new paradigm for SNARK design

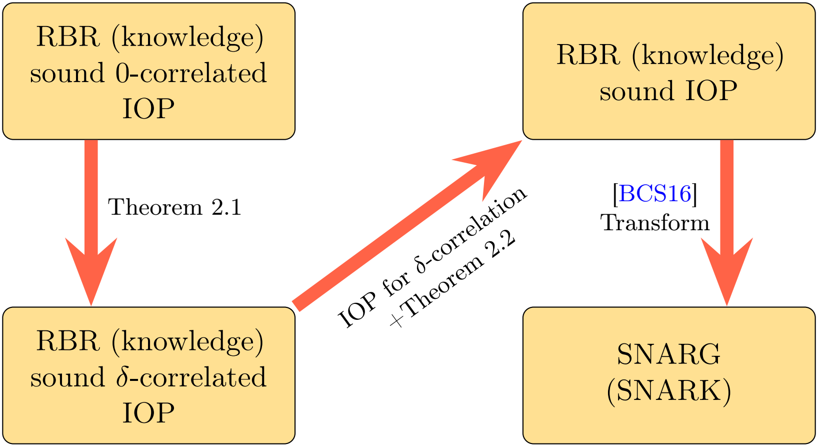

Our Results: "Plonk-like" δ-Correlated IOP

- We present a δ-correlated IOP called OPlonky
  - Captures "Plonk-like" protocols which use FRI as a sub-routine

Our Results: "Plonk-like" δ-Correlated IOP

- We present a δ-correlated IOP called OPlonky
  - Captures "Plonk-like" protocols which use FRI as a sub-routine
- We prove RBR soundness of OPlonky
  - Captures RBR soundness of [Pol, KPV22] and other Plonk-like protocols

Our Results: "Plonk-like" δ-Correlated IOP

- We present a δ-correlated IOP called OPlonky
  - Captures "Plonk-like" protocols which use FRI as a sub-routine
- We prove RBR soundness of OPlonky
  - Captures RBR soundness of [Pol, KPV22] and other Plonk-like protocols
- Our results can also be used to prove RBR soundness of ethSTARK and RISC Zero [Tea23]
  - ethSTARK has since independently been proven to be RBR sound [Sta23]

Remainder of the Talk

- Full FRI Protocol Overview
- RBR Soundness of FRI

Phase 1: Folding Phase

Phase 1: Folding Phase

Prover P

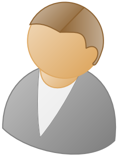

Verifier V

Phase 1: Folding Phase

$$G_0 \colon L_0 \to \mathbb{F}$$

Phase 1: Folding Phase

$G_0 \colon L_0 \to \mathbb{F}$

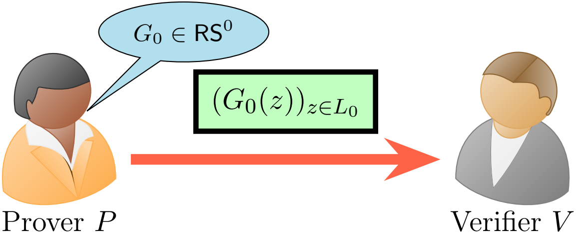

Phase 1: Folding Phase

$G_0 \colon L_0 \to \mathbb{F}$

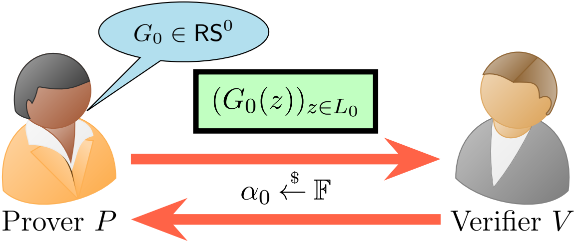

Phase 1: Folding Phase

$G_0 : L_0 \to \mathbb{F}$

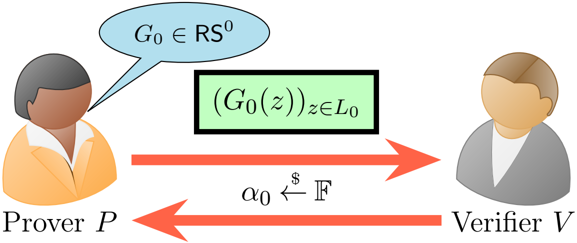

"Fold" G_0 into G_1

Phase 1: Folding Phase

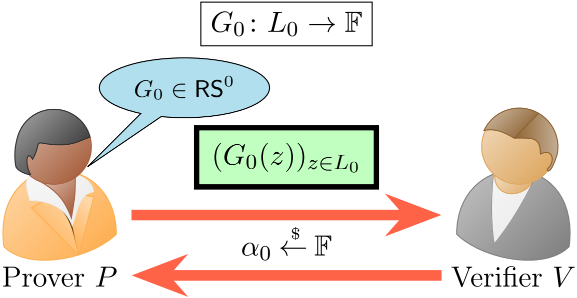

"Fold" G_0 into G_1

1. $L_1 := \{z^2 \colon z \in L_0\}$, $d_1 := d_0/2$

The FRI Protocol

Phase 1: Folding Phase

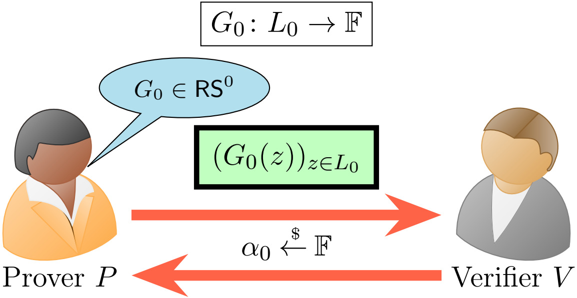

"Fold" G_0 into G_1

1. $L_1 := \{z^2 \colon z \in L_0\}$, $d_1 := d_0/2$
2. Define $G_1 \colon L_1 \to \mathbb{F}$ as random α_0-linear combo of "left" and "right" halves of G_0

Phase 1: Folding Phase

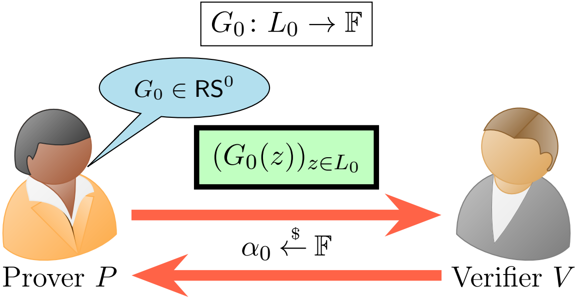

"Fold" G_0 into G_1

1. $L_1 := \{z^2 \colon z \in L_0\}$, $d_1 := d_0/2$
2. Define $G_1 \colon L_1 \to \mathbb{F}$ as random α_0-linear combo of "left" and "right" halves of G_0
3. Recurse above with G_1 and $\mathsf{RS}^1 = \mathsf{RS}[\mathbb{F}, L_1, d_1]$

Phase 2: Query Phase

$\log(d_0) = k$ rounds of folding

Prover P

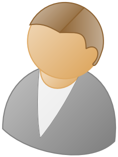

Verifier V

Phase 2: Query Phase

$\log(d_0) = k$ rounds of folding

Prover P

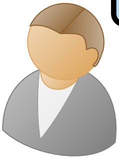

Verifier V

Has oracles $G_0, \ldots, G_{k-1}$

Phase 2: Query Phase

$\log(d_0) = k$ rounds of folding

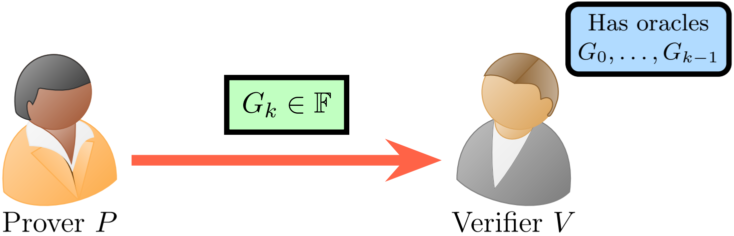

Phase 2: Query Phase

$\log(d_0) = k$ rounds of folding

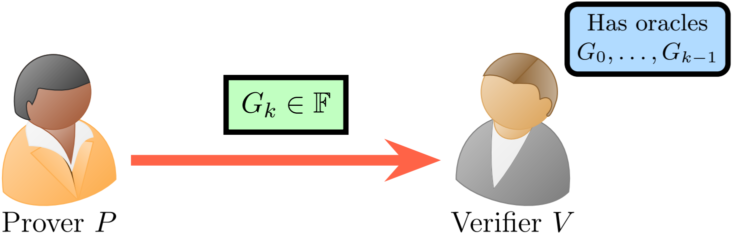

Check Consistency

Phase 2: Query Phase

$\log(d_0) = k$ rounds of folding

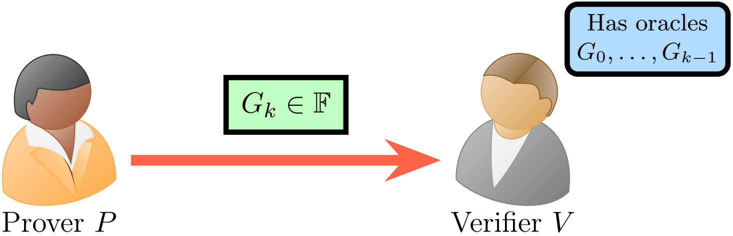

Check Consistency

1. V queries each oracle at 2 positions

The FRI Protocol

Phase 2: Query Phase

$\log(d_0) = k$ rounds of folding

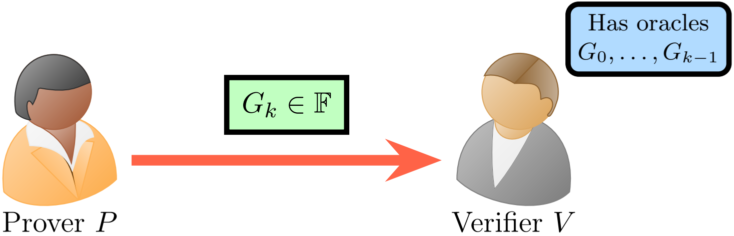

Check Consistency

1. V queries each oracle at 2 positions
2. V checks consistency of G_{i-1} and G_i for all $i \in [k]$

Phase 2: Query Phase

$\log(d_0) = k$ rounds of folding

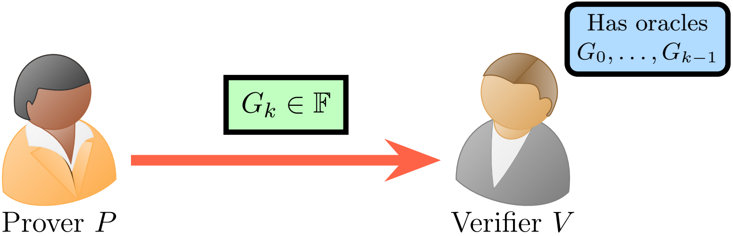

Check Consistency

1. V queries each oracle at 2 positions
2. V checks consistency of G_{i-1} and G_i for all $i \in [k]$
3. V repeats this process ℓ times

RBR Soundness of Folding Phase

RBR Soundness of FRI

RBR Soundness of Folding Phase

Protocol starts doomed

$G_0 \colon L_0 \to \mathbb{F}$ is δ-far from RS^0

RBR Soundness of Folding Phase

$G_0 \colon L_0 \to \mathbb{F}$ is δ-far from RS^0

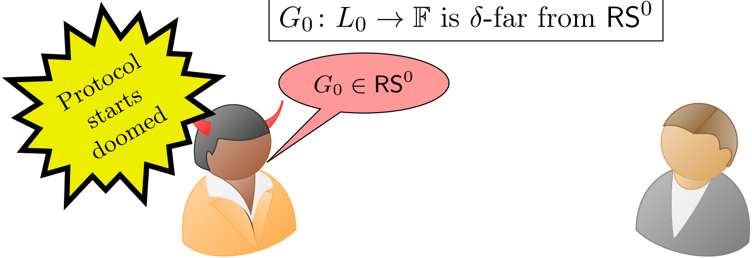

Prover P^*

Verifier V

RBR Soundness of Folding Phase

$G_0 \colon L_0 \to \mathbb{F}$ is δ-far from RS^0

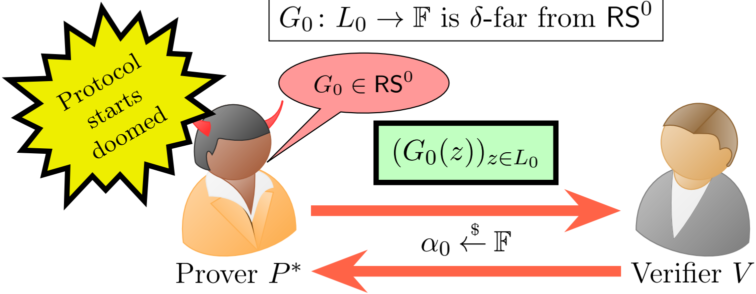

RBR Soundness of Folding Phase

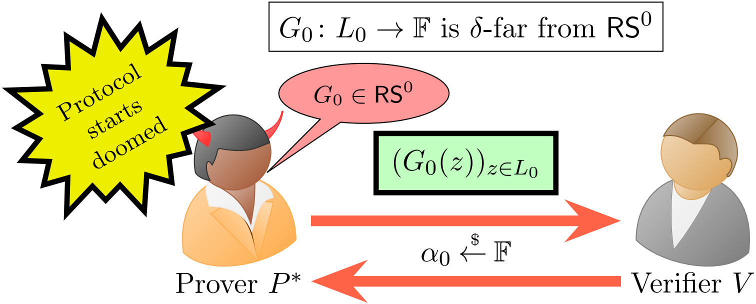

- If G_1 is δ-close to RS^1, then P^* can behave honestly and fool V!

RBR Soundness of Folding Phase

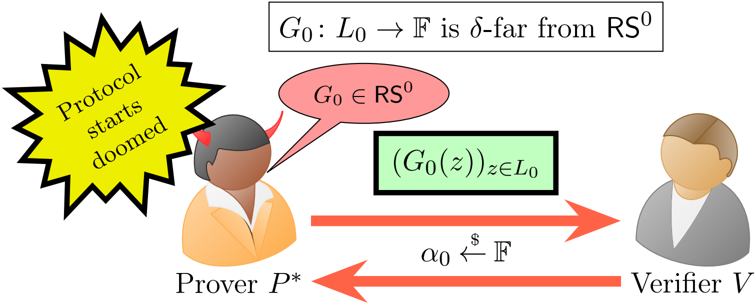

- If G_1 is δ-close to RS^1, then P^* can behave honestly and fool V!
- By [BCI+20]:

$$\Pr_{\alpha_0}[G_1 \text{ is } \delta\text{-close}] \leqslant \frac{(m+1/2)^7 |L_0|^2}{3\rho^{3/2}|\mathbb{F}|}.$$

RBR Soundness of FRI

RBR Soundness of Folding Phase

$G_{i-1} \colon L_{i-1} \to \mathbb{F}$ is δ-far from RS^{i-1}

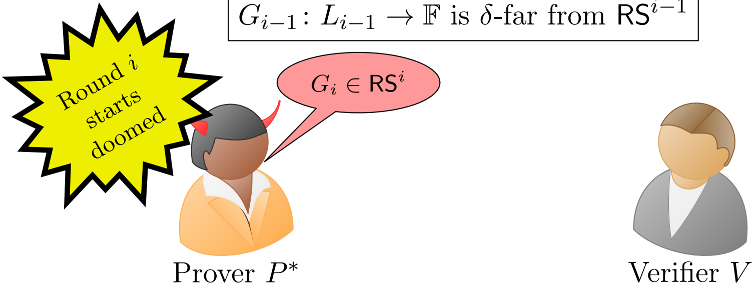

RBR Soundness of Folding Phase

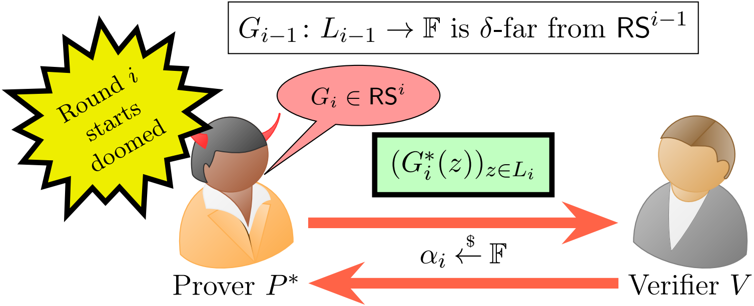

RBR Soundness of Folding Phase

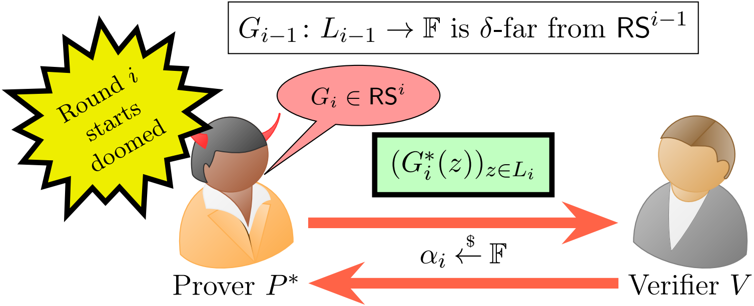

Protocol is doomed iff

RBR Soundness of Folding Phase

$G_{i-1} \colon L_{i-1} \to \mathbb{F}$ is δ-far from RS^{i-1}

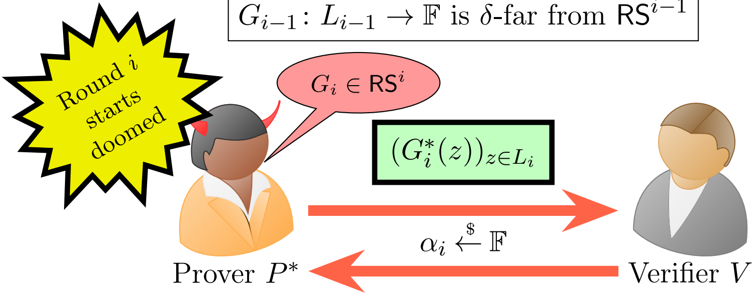

Protocol is doomed iff

1. G_i^* is not a correct folding of G_{i-1}; or

RBR Soundness of FRI

RBR Soundness of Folding Phase

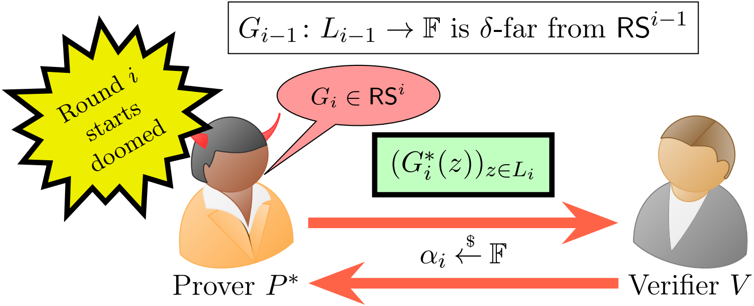

Protocol is doomed iff

1. G_i^* is not a correct folding of G_{i-1}; or
2. G_{i+1} (computed from honest G_i) is δ-far.

RBR Soundness of Folding Phase

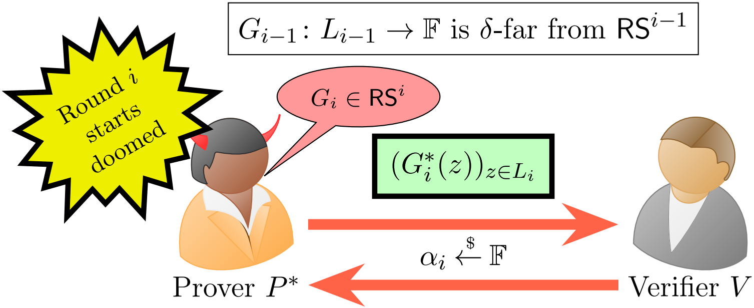

By same argument [BCI+20]:

$$\Pr_{\alpha_i}[G_i^*, \alpha_i \text{ is not doomed}] \leqslant \frac{(m+1/2)^7 |L_0|^2}{3\rho^{3/2}|\mathbb{F}|}.$$

RBR Soundness of Query Phase

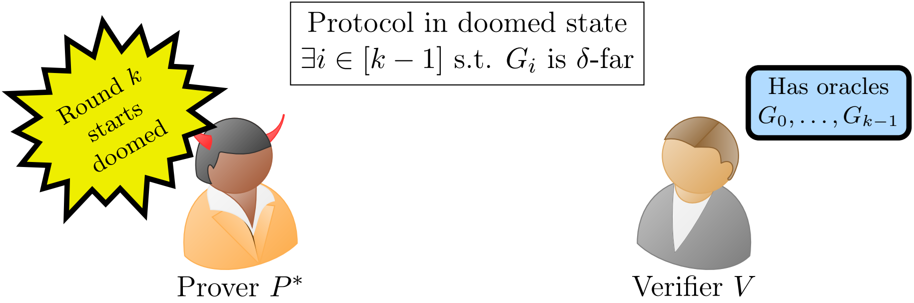

RBR Soundness of Query Phase

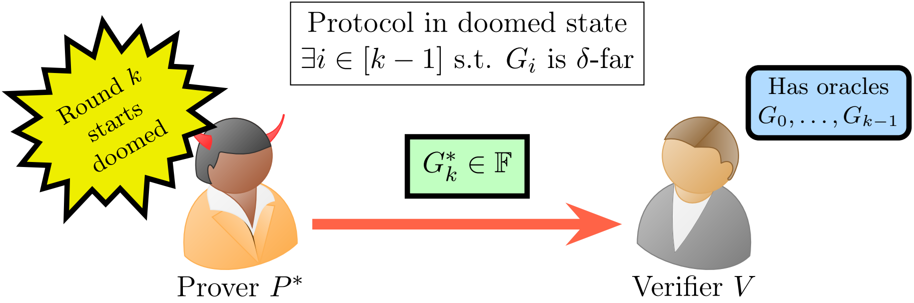

RBR Soundness of Query Phase

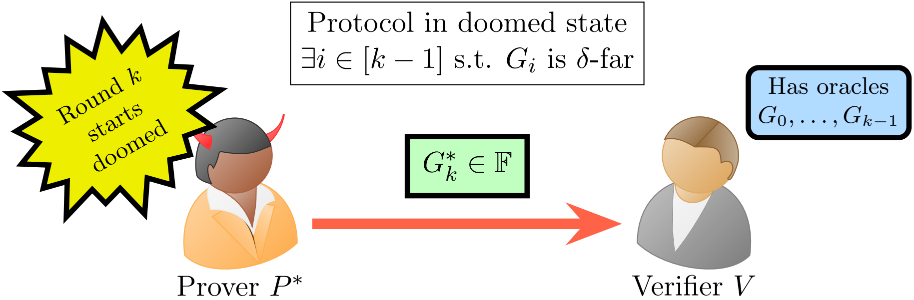

- Protocol is not doomed iff **all** V checks pass

RBR Soundness of Query Phase

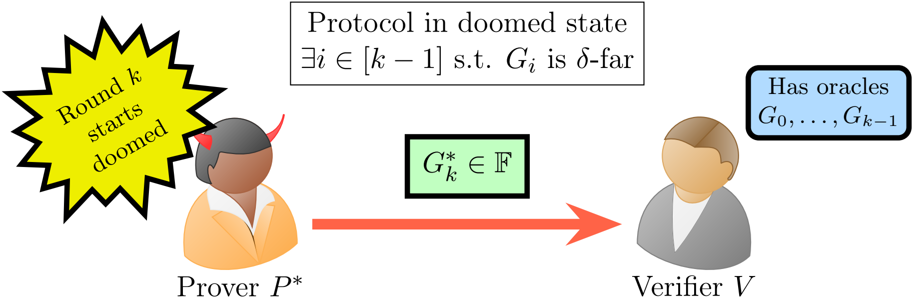

- Protocol is not doomed iff **all** V checks pass
- [BBH+18, BCI+20]: if $\exists i \in [k-1]$ such that G_i is δ-far, then

$$\Pr[\text{not doomed}] = \Pr[\text{all } V \text{ checks pass}] \leqslant (1-\delta)^{\ell}$$

Summary

Our Results: Bird's Eye View

1. Prove FS security of the **FRI Protocol** [BBH+18] and the **batched FRI Protocol**
   - Fills security gaps in [CMS19, COS20, KPV22]
2. Introduce **δ-Correlated IOPs** and prove their FS security
   - Intuitively, these are protocols that use FRI as a sub-routine
3. Formulate a δ-Correlated IOP which captures many "Plonk-like" protocols and prove their FS security
   - Captures Plonky2 [Pol], Redshift [KPV22], RISC Zero [Tea23]
   - ethSTARK [Sta23] and DEEP-FRI [BGK+20] also fit in this framework

Thank You!

Full version

https://ia.cr/2023/1071

References I

[Bar01] Boaz Barak. How to go beyond the black-box simulation barrier. In *FOCS*, pages 106–115. IEEE Computer Society, 2001.

[BBH+18] Eli Ben-Sasson, Iddo Bentov, Yinon Horesh, and Michael Riabzev. Fast reed-solomon interactive oracle proofs of proximity. In *ICALP*, volume 107 of *LIPIcs*, 14:1–14:17. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2018.

[BCG+18] Jonathan Bootle, Andrea Cerulli, Jens Groth, Sune Jakobsen, and Mary Maller. Arya: nearly linear-time zero-knowledge proofs for correct program execution. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 595–626. Springer, 2018.

[BCI+20] Eli Ben-Sasson, Dan Carmon, Yuval Ishai, Swastik Kopparty, and Shubhangi Saraf. Proximity gaps for reed-solomon codes. Cryptology ePrint Archive, Paper 2020/654, 2020. URL: `https://eprint.iacr.org/2020/654`. Full version of the same work published at FOCS 2020. DOI: `https://doi.org/10.1109/FOCS46700.2020.00088`.

[BCS16] Eli Ben-Sasson, Alessandro Chiesa, and Nicholas Spooner. Interactive oracle proofs. In *TCC (B2)*, volume 9986 of *Lecture Notes in Computer Science*, pages 31–60, 2016.

References II

[BDG+13] Nir Bitansky, Dana Dachman-Soled, Sanjam Garg, Abhishek Jain, Yael Tauman Kalai, Adriana López-Alt, and Daniel Wichs. Why "fiat-shamir for proofs" lacks a proof. In *TCC*, volume 7785 of *Lecture Notes in Computer Science*, pages 182–201. Springer, 2013.

[BEG+94] Manuel Blum, Will Evans, Peter Gemmell, Sampath Kannan, and Moni Naor. Checking the correctness of memories. *Algorithmica*, 12:225–244, 1994.

[BGK+20] Eli Ben-Sasson, Lior Goldberg, Swastik Kopparty, and Shubhangi Saraf. DEEP-FRI: sampling outside the box improves soundness. In *ITCS*, volume 151 of *LIPIcs*, 5:1–5:32. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020.

[CCH+19] Ran Canetti, Yilei Chen, Justin Holmgren, Alex Lombardi, Guy N. Rothblum, Ron D. Rothblum, and Daniel Wichs. Fiat-shamir: from practice to theory. In *STOC*, pages 1082–1090. ACM, 2019.

[CMS19] Alessandro Chiesa, Peter Manohar, and Nicholas Spooner. Succinct arguments in the quantum random oracle model. In *TCC (2)*, volume 11892 of *Lecture Notes in Computer Science*, pages 1–29. Springer, 2019.

References III

[COS20] Alessandro Chiesa, Dev Ojha, and Nicholas Spooner. Fractal: post-quantum and transparent recursive proofs from holography. In *EUROCRYPT (1)*, volume 12105 of *Lecture Notes in Computer Science*, pages 769–793. Springer, 2020.

[Dus] Dusk Network. Plonkup. https://github.com/dusk-network/plonkup. Accessed May 24, 2023.

[GK03] Shafi Goldwasser and Yael Tauman Kalai. On the (in)security of the fiat-shamir paradigm. In *FOCS*, pages 102–113. IEEE Computer Society, 2003.

[GWC19] Ariel Gabizon, Zachary J. Williamson, and Oana Ciobotaru. PLONK: permutations over lagrange-bases for oecumenical noninteractive arguments of knowledge. *IACR Cryptol. ePrint Arch.*:953, 2019.

[KPV22] Assimakis A. Kattis, Konstantin Panarin, and Alexander Vlasov. Redshift: transparent snarks from list polynomial commitments. In *CCS*, pages 1725–1737. ACM, 2022.

[Lip89] Richard J Lipton. *Fingerprinting sets*. Princeton University, Department of Computer Science, 1989.

[Lip90] Richard J Lipton. Efficient checking of computations. In *Annual Symposium on Theoretical Aspects of Computer Science*, pages 207–215. Springer, 1990.

References IV

[Mat] Matter Labs. Zksync 2.0: hello ethereum! https://blog.matter-labs.io/zksync-2-0-hello-ethereum-ca48588de179. Accessed May 24, 2023.

[Min] Mina. Mina book: background on plonk. https://o1-labs.github.io/proof-systems/plonk/overview.html. Accessed May 24, 2023.

[nil] =nil; Foundation. Circuit definition library for =nil; foundation's cryptography suite. https://github.com/NilFoundation/zkllvm-blueprint. Accessed May 24, 2023.

[Pol] Polygon Zero Team. Plonky2: fast recursive arguments with plonk and fri. URL: https://github.com/mir-protocol/plonky2/tree/main/plonky2. https://github.com/mir-protocol/plonky2/tree/main/plonky2.

[Sta23] StarkWare. Ethstark documentation v1.2. Cryptology ePrint Archive, Paper 2021/582, 2023. URL: https://eprint.iacr.org/2021/582. https://eprint.iacr.org/2021/582.

[Suc] Succinct Labs. Gnark-plonky2-verifier. https://github.com/succinctlabs/gnark-plonky2-verifier. Accessed May 24, 2023.

References V

[Tea23] RISC Zero Team. Risc zero's proof system for a zkvm. 2023. URL: https://github.com/risc0/risc0. Github repository.

[ZGK+18] Yupeng Zhang, Daniel Genkin, Jonathan Katz, Dimitrios Papadopoulos, and Charalampos Papamanthou. Vram: faster verifiable ram with program-independent preprocessing. In *2018 IEEE Symposium on Security and Privacy (SP)*, pages 908–925. IEEE, 2018.