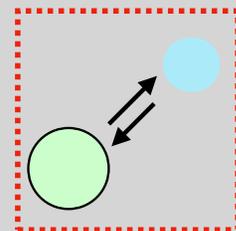


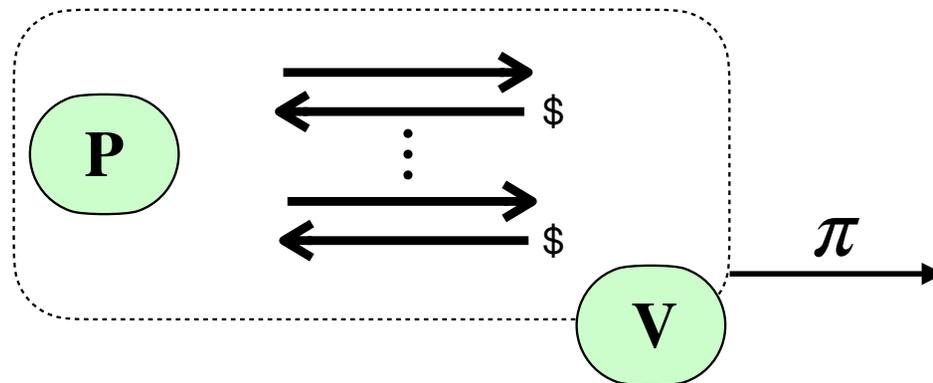
A Fiat–Shamir Transformation from Duplex Sponges

Alessandro Chiesa (EPFL), Michele Orrù (CNRS)



The Fiat-Shamir Transformation

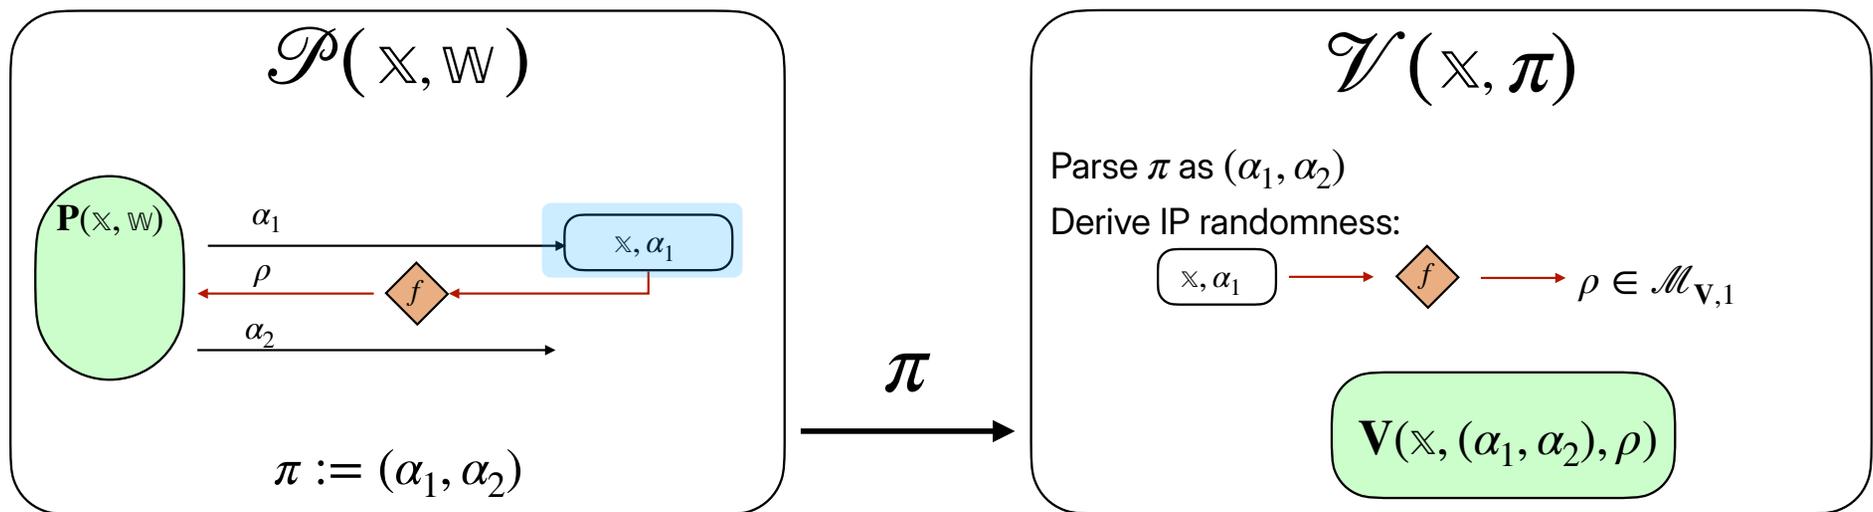
A public-coin interactive proof between a prover \mathbf{P} and a verifier \mathbf{V} .



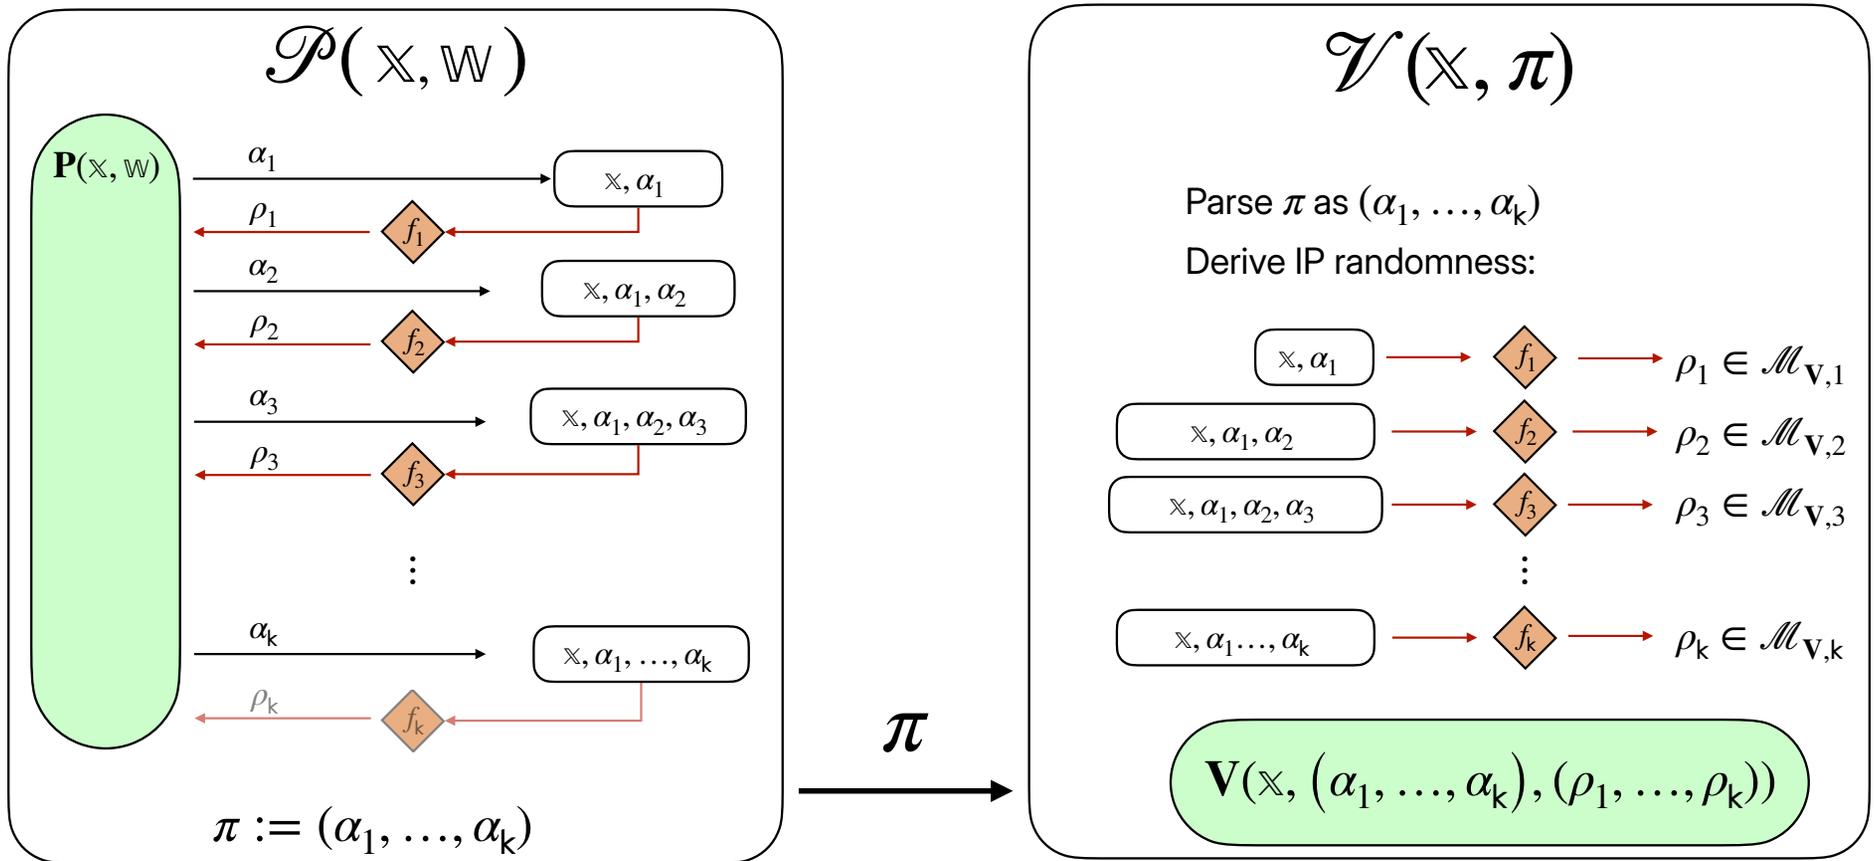
A hash function.



Fiat-Shamir is an umbrella term for multiple transformations.



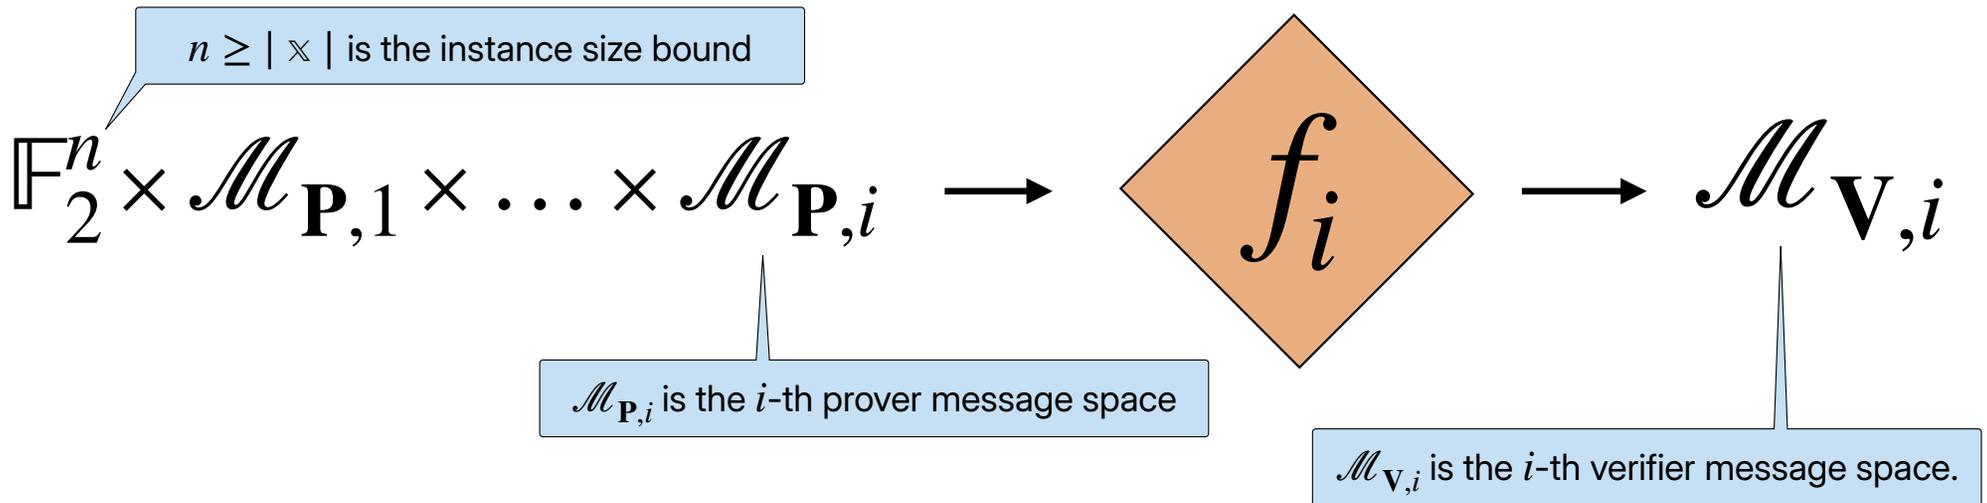
The canonical Fiat-Shamir transformation for a Σ -protocol.



The canonical Fiat–Shamir transformation for **multi-round** protocols.

What Kind of Random Oracles?

The canonical Fiat–Shamir transformation uses k **random oracles**.



Quadratic Blowup

The i -th query is contained in the $(i + 1)$ -th query. The overall query size is

$$k \cdot \text{len}(\times) + \sum_{i \in [k]} (k - i + 1) \cdot \text{len}(\alpha_i) = \Omega\left(\frac{k(k-1)}{2}\right)$$

Limitations

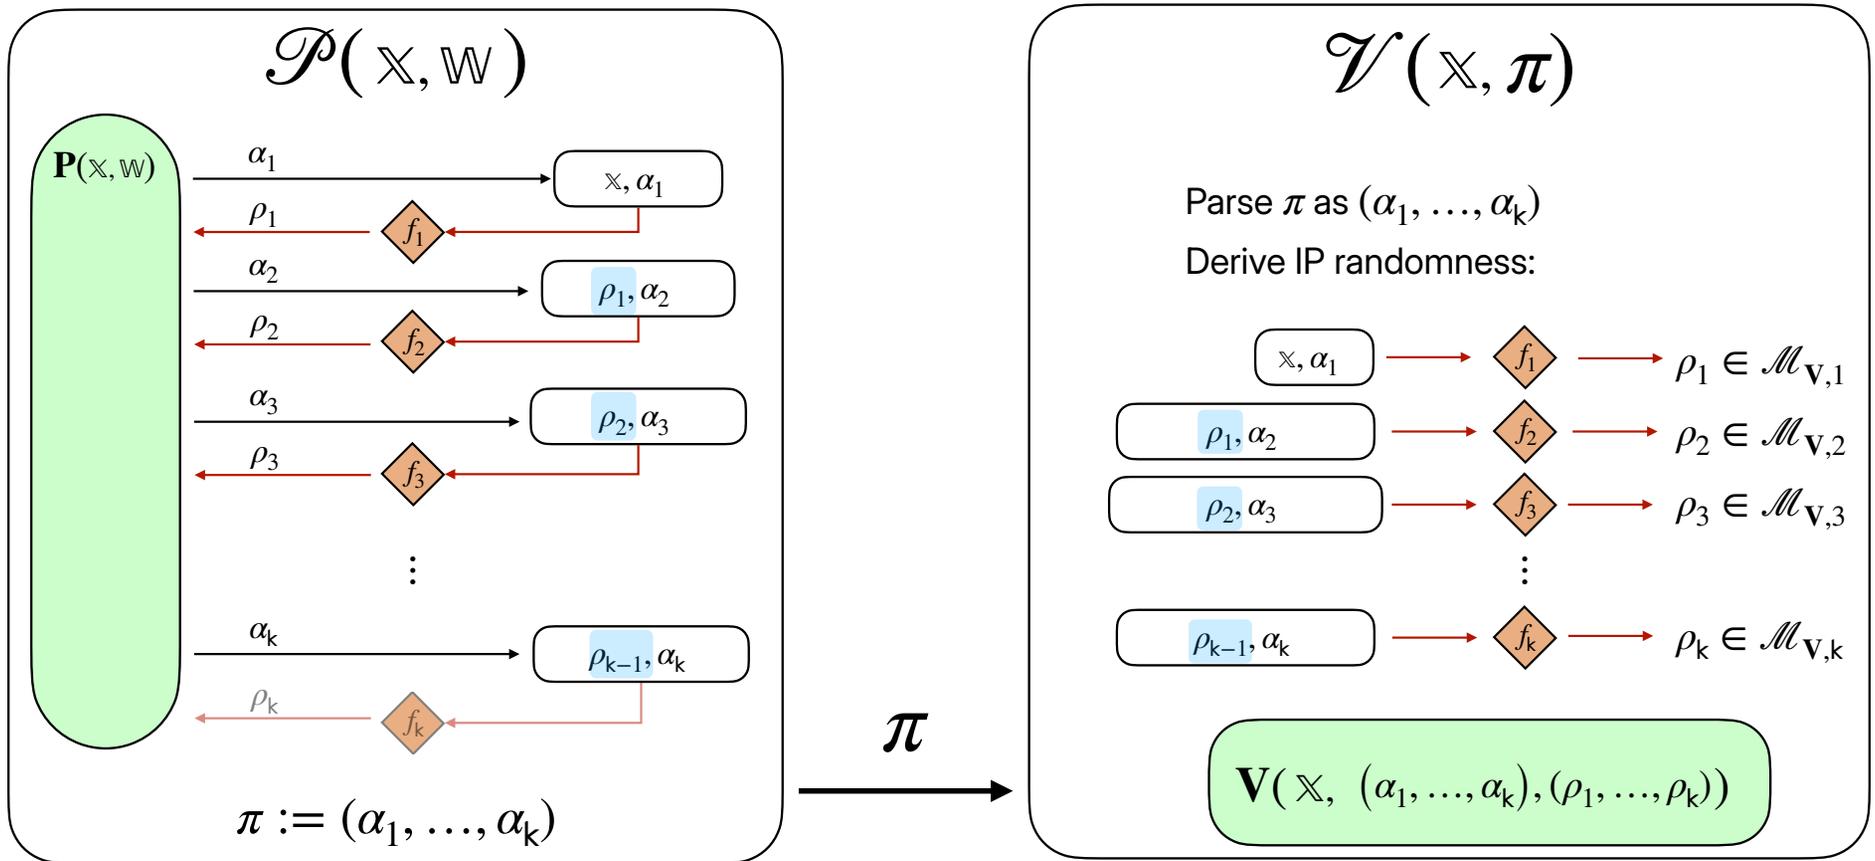
of the canonical Fiat–Shamir transformation.

Design Mismatch

Cryptographic hash functions rely on *blocks of fixed length*.

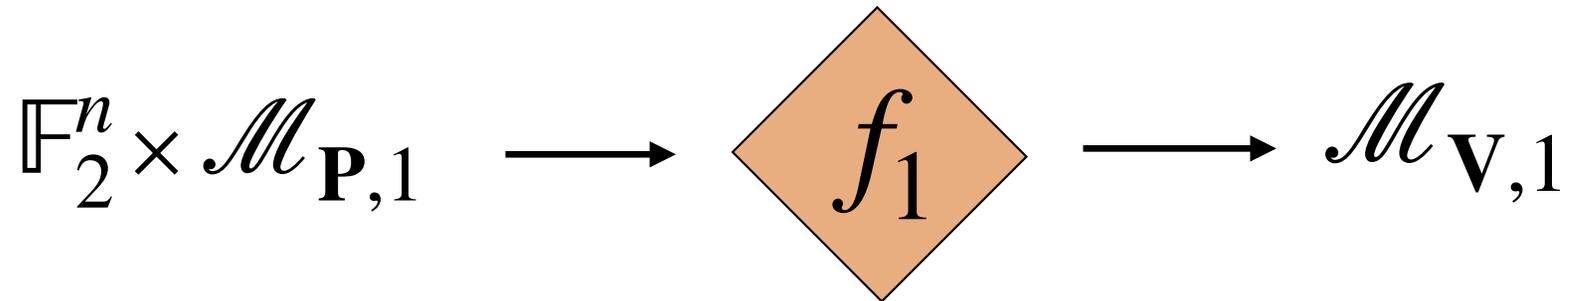
In canonical variant: $\mathbb{F}_2^n \times \mathcal{M}_{\mathbf{P},1} \times \dots \times \mathcal{M}_{\mathbf{P},i} \rightarrow \mathcal{M}_{\mathbf{V},i}$

In real world: $\Sigma^m \rightarrow \Sigma^n$.

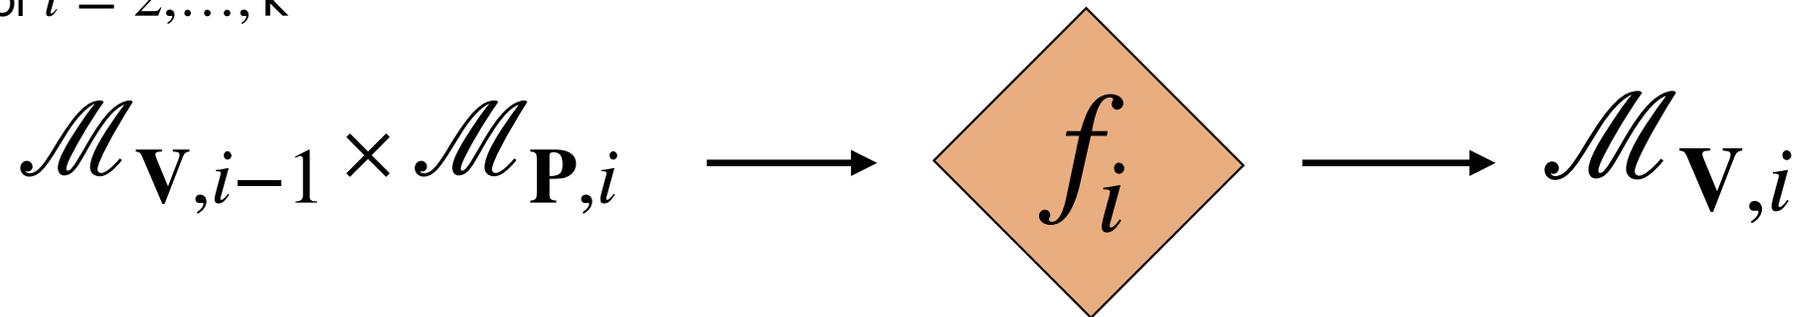


The hash-chain Fiat-Shamir transformation for multi-round protocols.

Oracles in the Hash-Chain Variant



For $i = 2, \dots, k$



Design Mismatch:

Oracles still depend on message spaces.

**Formally-studied
Fiat-Shamir transformations
are not used in practice**

Fiat-Shamir

Oracle $p \Rightarrow$ Real-world p

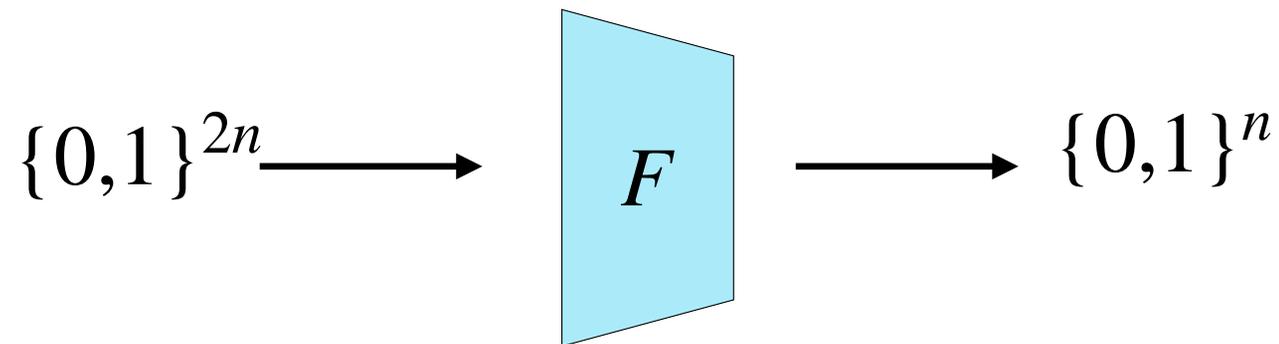
Fiat-Shamir

statistical soundness of IP
+
correlation intractability

Negative results:
[Bar01, GK03, BBHMR19, KRS25]

[IKRR16, CCRR18, HL18, CCHLRRW19, PS19,
LVW19, GJJM19, BFJKS19, LNPT19, CKU20, LV20,
BKM20, JKKZ20, CLMQ20, LNPY20, JJ21, HLR21,
CJJ21a, CJJ21b, LV22, HJKS22, GLS22,
BCHKLPR22, KLV23, CGJJZ23, DJJ24, IL25, ...]

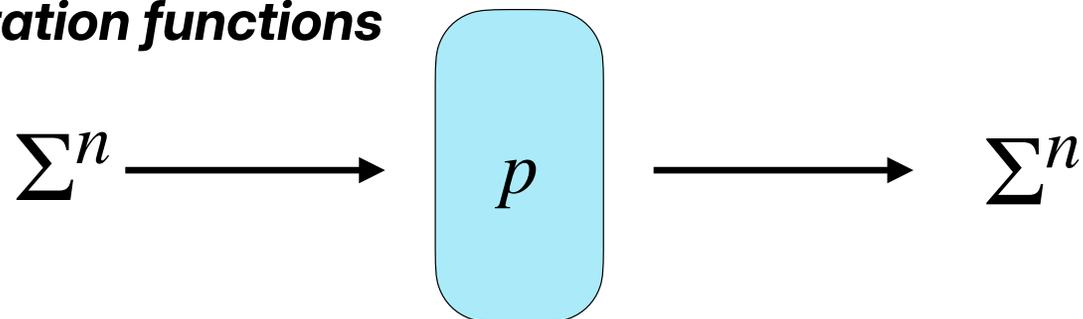
compression functions



ex: SHA-256, BLAKE2

In practice:

permutation functions

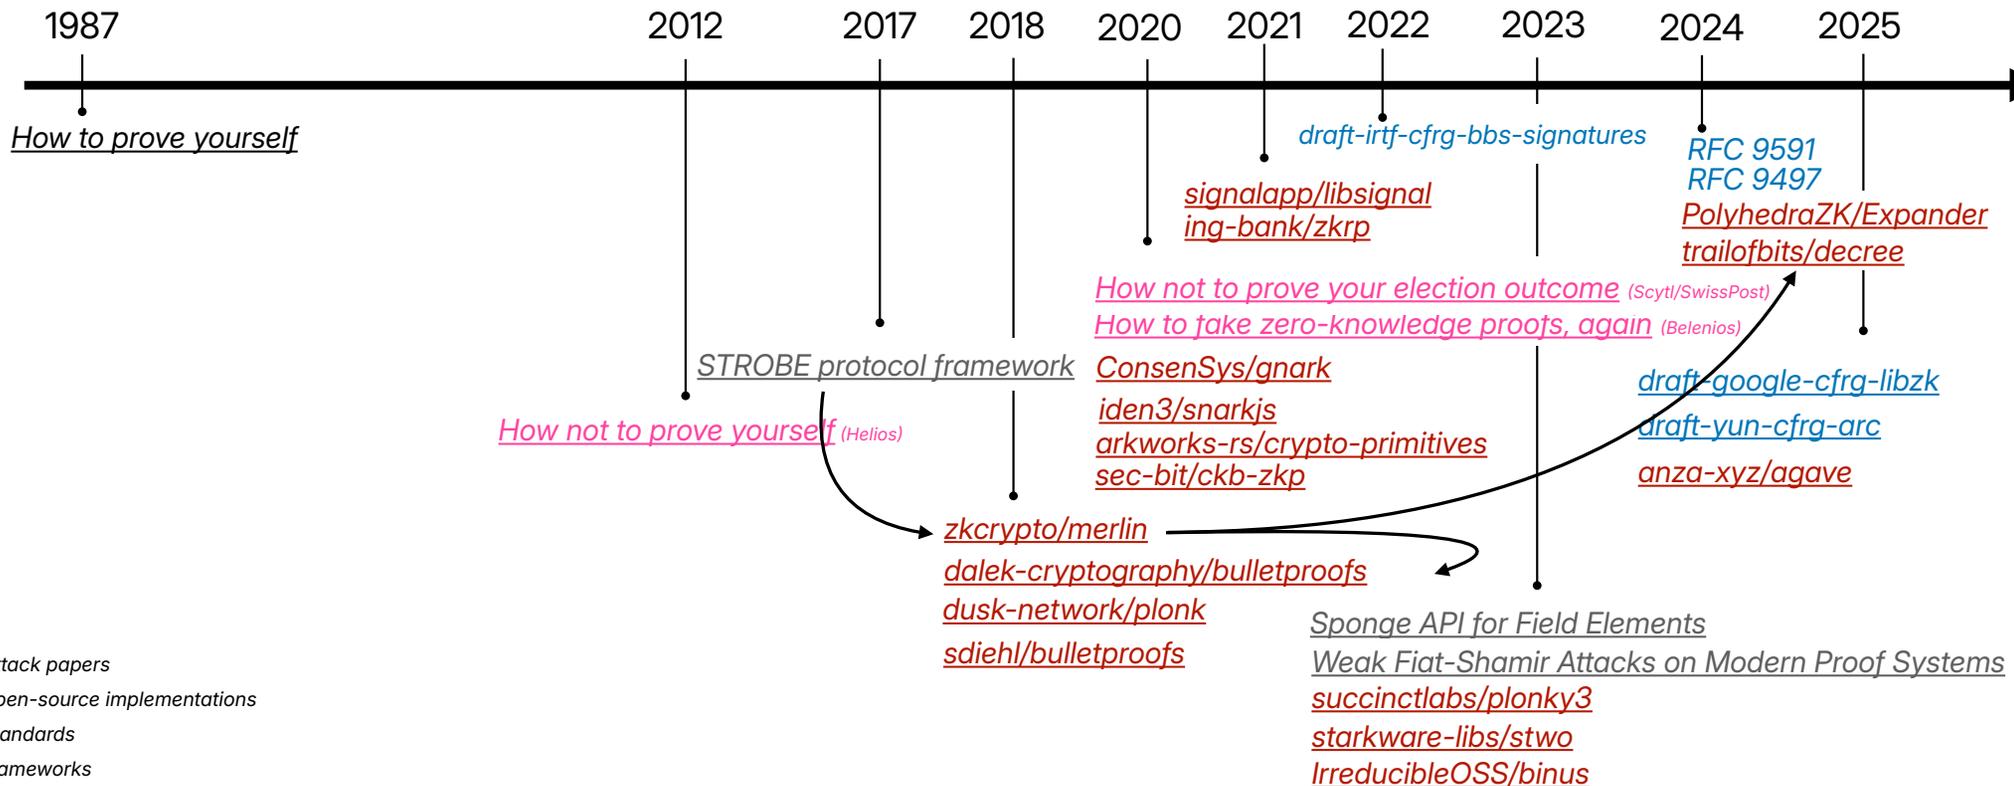


ex: Keccak-f, Poseidon

Unclear what are the security guarantees of modes of operation on top of these.

So... what do people do in practice?

Implementors have been left developing heuristics for themselves.



There's no rigorous analysis of any of these Fiat-Shamir transformations.

Our Contribution

A Fiat–Shamir transformation from ideal permutations



**Formal analysis
with precise security bounds.**

A Fiat–Shamir Transformation From Duplex Sponges

Alessandro Chiesa Michele Orrù
alessandro.chiesa@epfl.ch m@orru.net
EPFL CNRS

Abstract

The Fiat–Shamir transformation underlies numerous non-interactive arguments, with variants that differ in important ways. This paper addresses a gap between variants analyzed by theoreticians and variants implemented (and deployed) by practitioners. Specifically, theoretical analyses typically assume parties have access to random oracles with sufficiently large input and output size, while cryptographic hash functions in practice have fixed input and output sizes (pushing practitioners towards other variants).

In this paper we propose and analyze a variant of the Fiat–Shamir transformation that is based on an ideal permutation of fixed size. The transformation relies on the popular duplex sponge paradigm, and minimizes the number of calls to the permutation (given the amount of information to absorb and to squeeze). Our variant closely models deployed variants of the Fiat–Shamir transformation, and our analysis provides concrete security bounds that can be used to set security parameters in practice.

We additionally contribute `spongefish`, an open-source Rust library implementing our Fiat–Shamir transformation. The library is interoperable across multiple cryptographic frameworks, and works with any choice of permutation. The library comes equipped with Keccak and Poseidon permutations, as well as several “codecs” for re-mapping prover and verifier messages to the permutation’s domain.

Keywords: Fiat–Shamir transformation; duplex sponge



Open-source implementation.

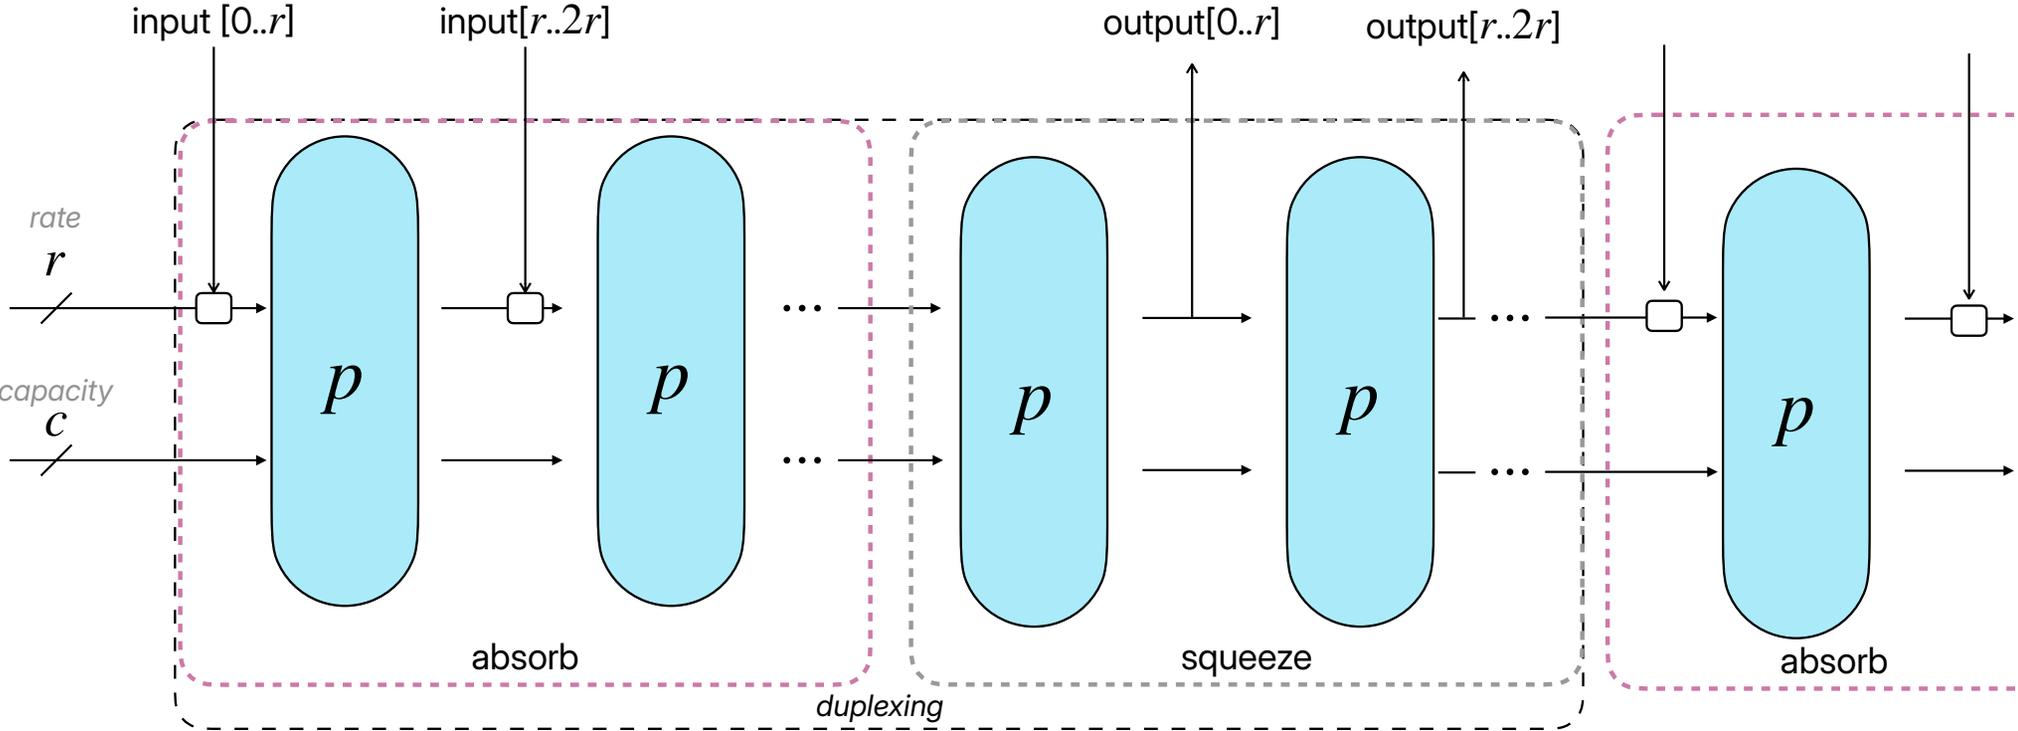
The screenshot shows the GitHub repository for 'spongefish'. At the top, it indicates the repository is public and has 2 branches and 1 tag. Below this, there is a search bar and buttons for 'Add file' and 'Code'. The main content is a list of recent commits:

Commit Message	Time
Attempt update lint-fmt.yml	3 weeks ago
Add latex rendering.	2 years ago
edits: clippy: add clippy constraints to other crates (#40)	3 weeks ago
edits: clippy: add clippy constraints to other crates (#40)	3 weeks ago
Add unit tests for blake3 Proof of Work (#53)	2 weeks ago
Simplify some map patterns (#55)	last week
Remove Cargo.lock	11 months ago
Remove unused anyhow dep (#56)	last week
Create LICENSE	2 years ago
Refactor using naming conventions of the academic paper.	3 weeks ago

Duplexing the sponge

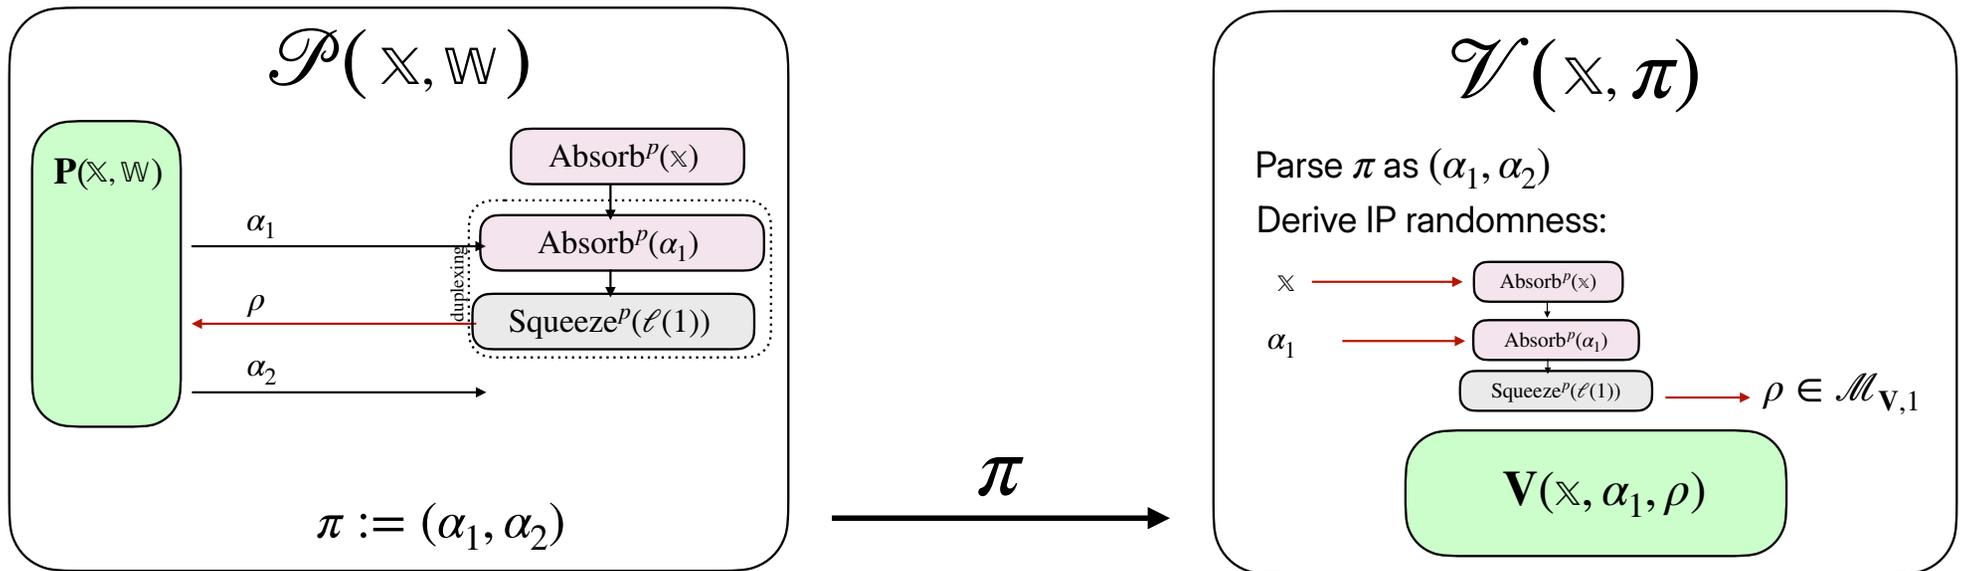
[BDPV07]
[BDPV11]

Let $p: \Sigma^{r+c} \rightarrow \Sigma^{r+c}$ be a permutation function.

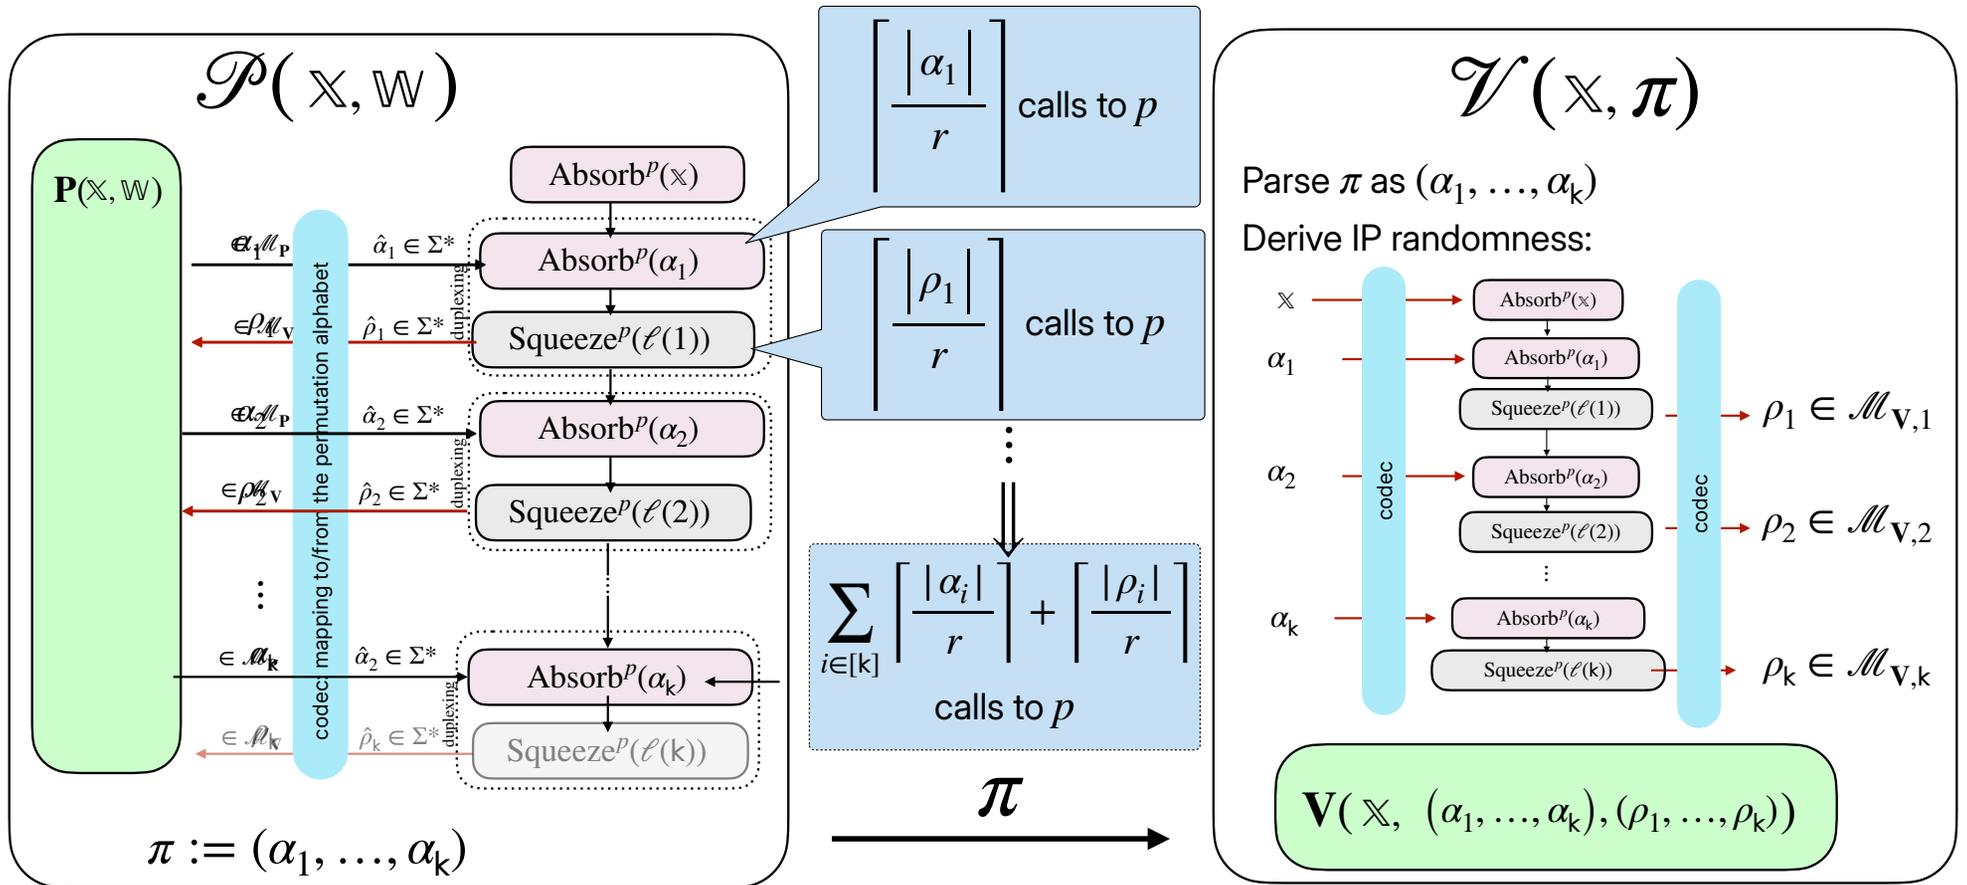


Absorb prover messages.
Squeeze verifier messages.

DSFS[IP] for Σ -protocols



DSFS[IP] for multi-round protocols



Proving security

In an oracle model, where $p: \Sigma^{r+c} \rightarrow \Sigma^{r+c}$ is an ideal permutation.

**(knowledge)
soundness**

[malicious prover]

$$\mathbb{X}, \pi \implies \mathbb{W}$$

$$\epsilon_{\text{srsnd}} + O\left(\frac{t^2}{|\Sigma|^c}\right)$$

zero-knowledge

[malicious verifier]

$$\mathcal{P}(\mathbb{X}, \mathbb{W}) \approx \mathcal{S}(\mathbb{X})$$

$$\epsilon_{\text{hvzk}} + O\left(\frac{t}{|\Sigma|^{\min(c, \delta)}}\right)$$

δ is the **privacy parameter**

where t is the number of queries by adversary to p .

Soundness

$$\mathbb{X}, \pi \Longrightarrow \mathbb{W}$$

(Knowledge) Soundness

$$\begin{array}{l}
 p \xleftarrow{\text{tr}} \mathcal{D}_\varepsilon \\
 (\mathbb{X}, \pi) \xleftarrow{\tilde{\mathcal{J}}^{p,p^{-1}}} \quad \longrightarrow \quad \tilde{\mathbf{P}}^{\text{SR}}
 \end{array}$$

- Soundness:

$$\varepsilon_{\text{snd}} = \Pr [\pi \text{ valid and } \mathbb{X} \notin \mathcal{L}]$$

- Straightline Knowledge Soundness:

$$\kappa = \Pr [\mathbb{W} \leftarrow \mathcal{E}(\mathbb{X}, \pi, \text{tr}) \text{ such that } \pi \text{ valid and } (\mathbb{X}, \mathbb{W}) \notin \mathcal{R}]$$

- Rewinding Knowledge Soundness:

$$\kappa_{\text{rw}} = \Pr [\mathbb{W} \leftarrow \mathcal{E}(\mathbb{X}, \pi, \text{tr}, \mathcal{P}) \text{ such that } \pi \text{ valid and } (\mathbb{X}, \mathbb{W}) \notin \mathcal{R}]$$

$$\varepsilon \leq \kappa \leq \kappa_{\text{rw}}$$

- State-Restoration Soundness

$$\varepsilon^{(\text{sr})} \leq \varepsilon$$

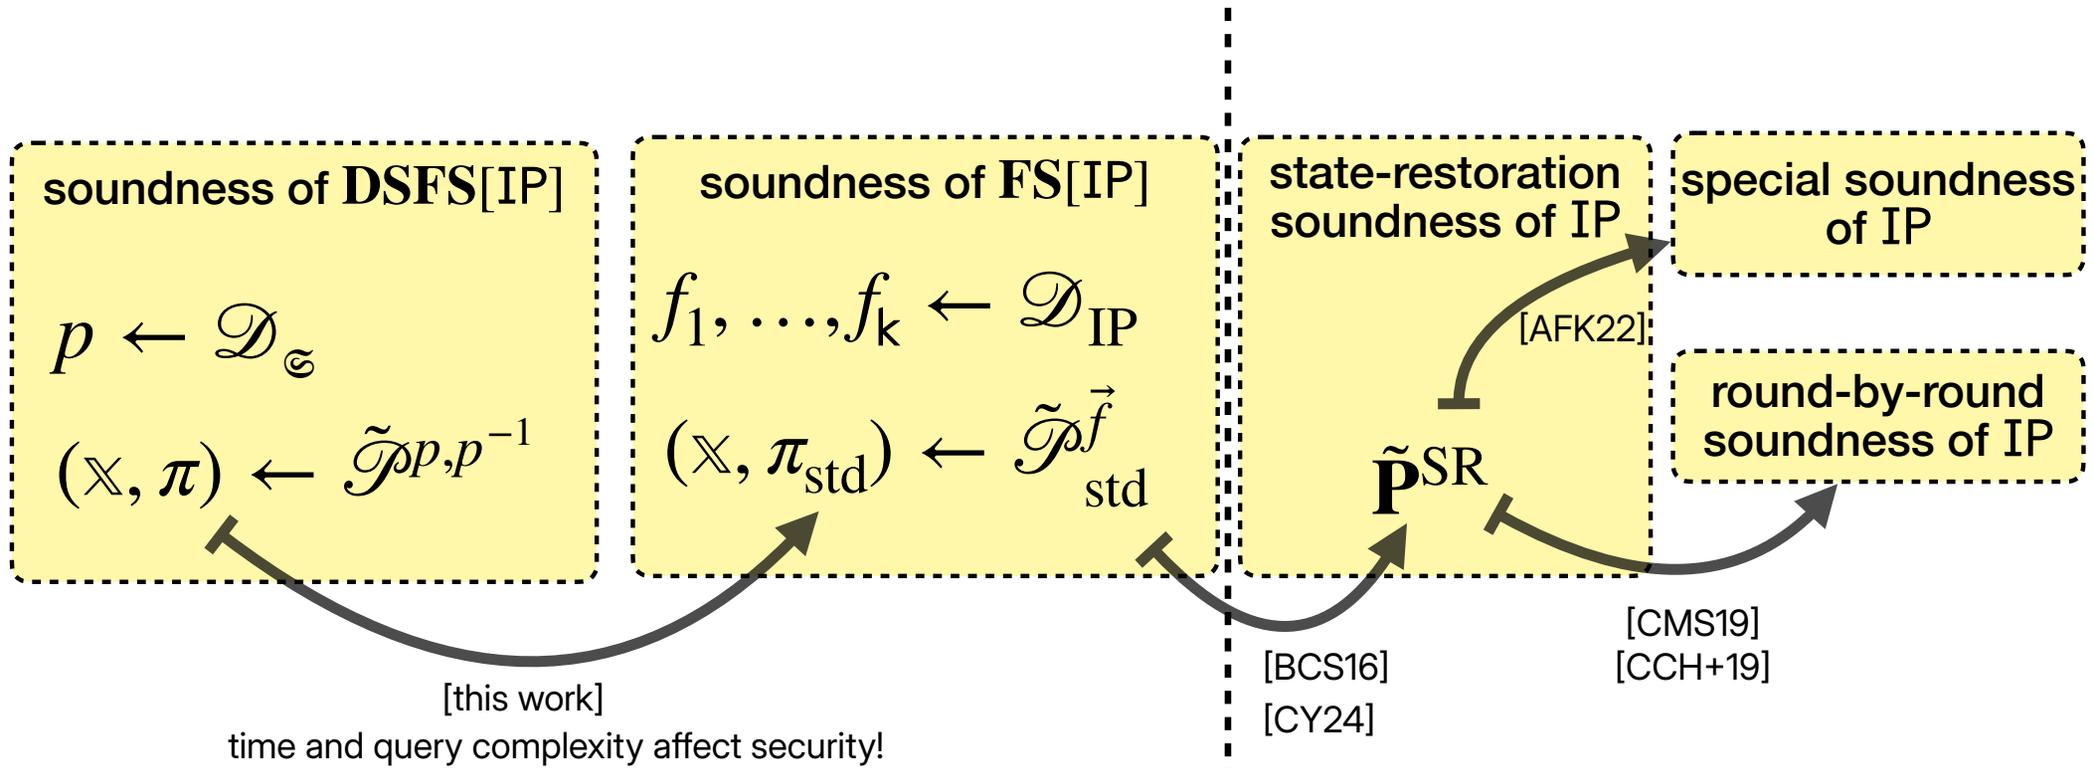
- State-Restoration Knowledge Soundness

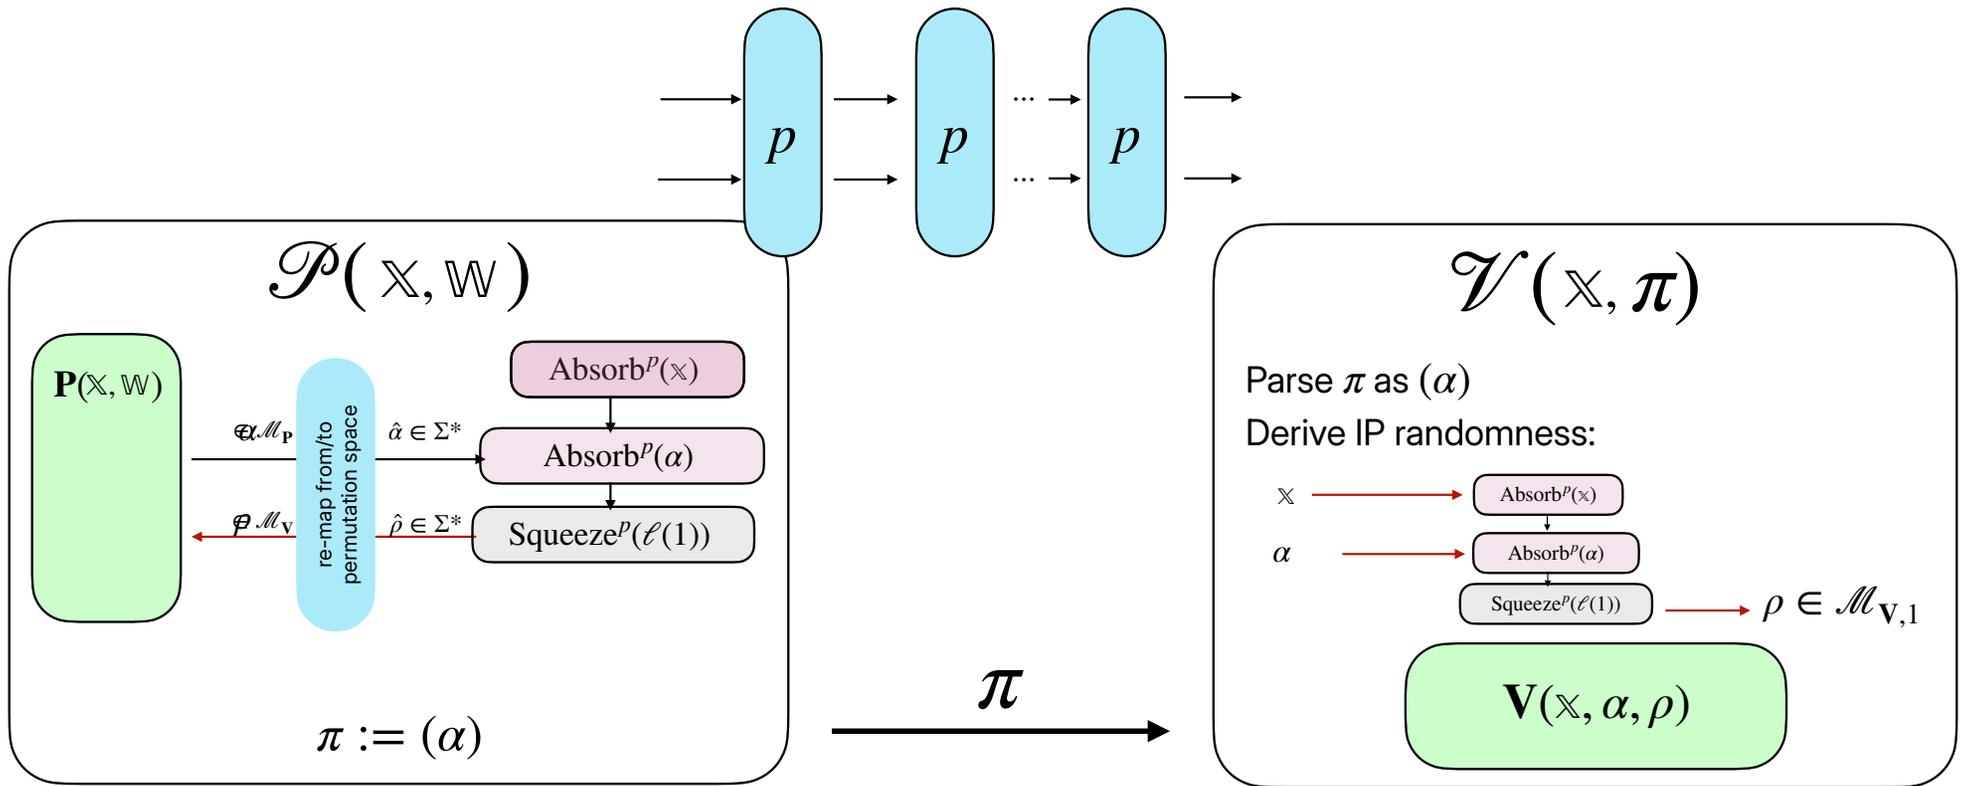
$$\kappa^{(\text{sr})} \leq \kappa$$

- State-Restoration Rewinding Knowledge Soundness

$$\kappa_{\text{rw}}^{(\text{sr})} \leq \kappa_{\text{rw}}$$

Proving soundness: strategy





Proving soundness: main lemma

We seek a reduction \mathcal{P}_{std} such that:

$$p \leftarrow \mathcal{D}_{\mathfrak{G}}$$

$$f_1, \dots, f_k \leftarrow \mathcal{D}_{\text{IP}}$$

$$(\mathbb{X}, \pi) \xleftarrow{\text{tr}} \tilde{\mathcal{P}}^{p, p^{-1}}$$

$$(\mathbb{X}, \pi_{\text{std}}) \xleftarrow{\text{tr}_{\text{std}}} \tilde{\mathcal{P}}_{\text{std}}^{\vec{f}}$$

Proving soundness: main lemma

Find ReduceP and ReduceT such that:

$$p \leftarrow \mathcal{D}_{\mathfrak{G}}$$

$$f_1, \dots, f_k \leftarrow \mathcal{D}_{\text{IP}}$$

$$(\mathbb{X}, \pi) \xleftarrow{\text{tr}} \tilde{\mathcal{P}}^{p, p^{-1}} \iff (\mathbb{X}, \pi_{\text{std}}) \xleftarrow{\text{tr}_{\text{std}}} \text{ReduceP}(\tilde{\mathcal{P}})^{\vec{f}}$$

$$b \leftarrow \mathcal{V}(\mathbb{X}, \pi)$$

$$\approx$$

$$b_{\text{std}} \leftarrow \mathcal{V}_{\text{std}}(\mathbb{X}, \pi_{\text{std}})$$

$$\text{tr}_{\text{std}} \leftarrow \text{ReduceT}(\text{tr})$$

$$\text{tr}_{\text{std}}$$

$$\text{Additive error: } \eta_{\star} = O\left(\frac{t^2}{|\Sigma|^c}\right).$$

Isn't this already known?

After all, we already had the sponge construction, and know how to build XOFs.

Indistinguishability

Let $\mathbf{f} \leftarrow \mathcal{D}_{\text{IP}}$ and $p \leftarrow \mathcal{D}_{\mathcal{C}}$. A construction \mathcal{C} is **indistinguishable** if:

$$\left\{ \begin{array}{c} b \leftarrow \mathcal{A}^{\mathbf{f}} \\ b \end{array} \right\} \approx \left\{ \begin{array}{c} b \leftarrow \mathcal{A}^{\mathcal{C}^p} \\ b \end{array} \right\}$$

The adversary $\mathcal{A}^{p,p^{-1}}$ can query the permutation function.

The Indifferentiability Framework

Let $\mathbf{f} \leftarrow \mathcal{D}_{\text{IP}}$ and $p \leftarrow \mathcal{D}_{\mathcal{E}}$. A construction \mathcal{C} is **indifferentiable** $\exists \mathcal{S}$:

$$\left\{ \begin{array}{c} b \leftarrow \mathcal{A}^{\mathbf{f}, \mathcal{S}^{\mathbf{f}}} \\ b \end{array} \right\} \approx \left\{ \begin{array}{c} b \leftarrow \mathcal{A}^{\mathcal{C}^p, p} \\ b \end{array} \right\}$$

Soundness can be proven with this property

The extractor \mathcal{E} requires the random oracle trace.

$$\varepsilon_{\text{ksnd}} = \Pr \left[w \leftarrow \mathcal{E}(x, \pi, \text{tr}) \text{ such that } \pi \text{ valid and } (x, w) \notin \mathcal{R} \right]$$

Our notion: double indiffereniability

Let $\mathbf{f} \leftarrow \mathcal{D}_{\text{IP}}$ and $p \leftarrow \mathcal{D}_{\mathfrak{C}}$. A construction \mathcal{C} is **doubly indiffereniiable** if $\exists \mathcal{S}, \mathcal{T}$:

$$\left\{ \begin{array}{l} \text{out} \xleftarrow{\text{tr}} \mathcal{A}^{\mathbf{f}, \mathcal{S}^{\mathbf{f}}} \\ \text{tr}, \text{out} \end{array} \right\} \approx \left\{ \begin{array}{l} \text{out} \xleftarrow{\text{tr}_{\diamond}} \mathcal{A}^{\mathcal{C}^p, p} \\ \mathcal{T}(\text{tr}_{\diamond}), \text{out} \end{array} \right\}$$

Our notion: double indiffereniability

Let $\mathbf{f} \leftarrow \mathcal{D}_{\text{IP}}$ and $p \leftarrow \mathcal{D}_{\mathfrak{G}}$. A construction \mathcal{C} is **doubly indiffereniiable** if $\exists \mathcal{S}, \mathcal{T}$:

$$\left\{ \begin{array}{l} \text{out} \xleftarrow{\text{tr}} \mathcal{A}^{\mathbf{f}, \mathcal{S}^{\mathbf{f}}} \\ \text{tr}, \text{out} \end{array} \right\} \approx \left\{ \begin{array}{l} \text{out} \xleftarrow{\text{tr}_{\diamond}} \mathcal{A}^{\mathcal{C}^p, p} \\ \mathcal{T}(\text{tr}_{\diamond}), \text{out} \end{array} \right\}$$

$$\kappa_{\text{DSFS}[\text{IP}]} = \Pr \left[b = 1 \wedge (x, \pi) \mid \begin{array}{l} p \leftarrow \mathcal{D}_{\mathfrak{G}} \\ (x, \pi) \xleftarrow{\text{tr}_{\diamond}} \mathcal{A}^p \\ \mathbb{W} \leftarrow \mathcal{E}_{\text{DSFS}}(x, \pi, \text{tr}_{\diamond}) \\ b \leftarrow \mathcal{V}^p(x, \pi) \end{array} \right] = \Pr \left[b = 1 \wedge (x, \pi) \mid \begin{array}{l} p \leftarrow \mathcal{D}_{\mathfrak{G}} \\ (x, \pi) \xleftarrow{\text{tr}_{\diamond}} \mathcal{A}^p \\ \text{tr} \leftarrow \mathcal{T}(\text{tr}_{\diamond}) \\ \mathbb{W} \leftarrow \mathcal{E}_{\text{IP}}(x, \pi, \text{tr}_{\diamond}) \\ b \leftarrow \mathcal{V}^p(x, \pi) \end{array} \right]$$

Our notion: double indiffereniability

Let $\mathbf{f} \leftarrow \mathcal{D}_{\text{IP}}$ and $p \leftarrow \mathcal{D}_{\mathcal{G}}$. A construction \mathcal{C} is **doubly indiffereniiable** if $\exists \mathcal{S}, \mathcal{T}$:

$$\left\{ \begin{array}{l} \text{out} \xleftarrow{\text{tr}} \mathcal{A}^{\mathbf{f}, \mathcal{S}^{\mathbf{f}}} \\ \text{tr}, \text{out} \end{array} \right\} \approx \left\{ \begin{array}{l} \text{out} \xleftarrow{\text{tr}_{\diamond}} \mathcal{A}^{\mathcal{C}^p, p} \\ \mathcal{T}(\text{tr}_{\diamond}), \text{out} \end{array} \right\}$$

$$\kappa_{\text{DSFS}[\text{IP}]} = \Pr \left[b = 1 \wedge (x, \pi) \mid \begin{array}{l} p \leftarrow \mathcal{D}_{\mathcal{G}} \\ (x, \pi) \leftarrow \mathcal{A}^p \\ \text{tr} \leftarrow \mathcal{T}(\text{tr}_{\diamond}) \\ \mathbb{W} \leftarrow \mathcal{E}_{\text{IP}}(x, \pi, \text{tr}_{\diamond}) \\ b \leftarrow \mathcal{V}^p(x, \pi) \end{array} \right] \leq \Pr \left[b = 1 \wedge (x, \pi) \mid \begin{array}{l} \mathbf{f} \leftarrow \mathcal{D}_{\text{IP}} \\ (x, \pi) \xleftarrow{\text{tr}} \mathcal{A}^{\mathcal{S}^{\mathbf{f}}} \\ \mathbb{W} \leftarrow \mathcal{E}_{\text{IP}}(x, \pi, \text{tr}) \\ b \leftarrow \mathcal{V}^{\mathcal{S}^{\mathbf{f}}}(x, \pi) \end{array} \right] + \mu_{\star}$$

Our notion: double indifferenciability

Let $\mathbf{f} \leftarrow \mathcal{D}_{\text{IP}}$ and $p \leftarrow \mathcal{D}_{\mathcal{E}}$. A construction \mathcal{C} is **doubly indifferenciabile** if $\exists \mathcal{S}, \mathcal{T}$:

$$\left\{ \begin{array}{c} \text{out} \xleftarrow{\text{tr}} \mathcal{A}^{\mathbf{f}, \mathcal{S}^{\mathbf{f}}} \\ \text{tr}, \text{out} \end{array} \right\} \approx \left\{ \begin{array}{c} \text{out} \xleftarrow{\text{tr}_{\diamond}} \mathcal{A}^{\mathcal{E}^p, p} \\ \mathcal{T}(\text{tr}_{\diamond}), \text{out} \end{array} \right\}$$

$$\kappa_{\text{DSFS}[\text{IP}]} \leq \Pr \left[b = 1 \wedge (x, \pi) \mid \begin{array}{c} \mathbf{f} \leftarrow \mathcal{D}_{\text{IP}} \\ (x, \pi) \xleftarrow{\text{tr}} \mathcal{A}^{\mathcal{S}^{\mathbf{f}}} \\ w \leftarrow \mathcal{E}_{\text{IP}}(x, \pi, \text{tr}) \\ b \leftarrow \mathcal{V}^{\mathcal{S}^{\mathbf{f}}}(x, \pi) \end{array} \right] + \mu_{\star} = \kappa_{\text{snd}} + \mu_{\star}$$

Zero-Knowledge

$$\mathcal{P}(x, w) \approx \mathcal{S}(x)$$

Zero-Knowledge

In the explicitly-programmable random oracle model.

$$p \leftarrow \mathcal{D}_{\mathfrak{S}}$$

$$(\mathbb{X}, \mathbb{W}) \leftarrow \mathcal{A}^{p, p^{-1}}$$

Real-world prover

$$\pi \leftarrow \mathcal{P}^p(\mathbb{X}, \mathbb{W})$$

$$b_0 \leftarrow \mathcal{A}^{p, p^{-1}}(\pi)$$

Ideal-world simulator

$$(\pi, \mu) \leftarrow \mathcal{S}^{p, p^{-1}}(\mathbb{X})$$

$$b_1 \leftarrow \mathcal{A}^{(p, p^{-1})[\mu]}(\pi)$$



Proving zero-knowledge: what about indifferentiability?

Indifferentiability is insufficient.

1. The simulator must reprogram the random oracle.

Real-world prover

$$\pi \leftarrow \mathcal{P}^p(\mathbb{X}, \mathbb{W})$$

$$b_0 \leftarrow \mathcal{A}^{p,p^{-1}}(\pi)$$

Ideal-world Simulator

$$(\pi, \mu) \leftarrow \mathcal{S}^{p,p^{-1}}(\mathbb{X})$$

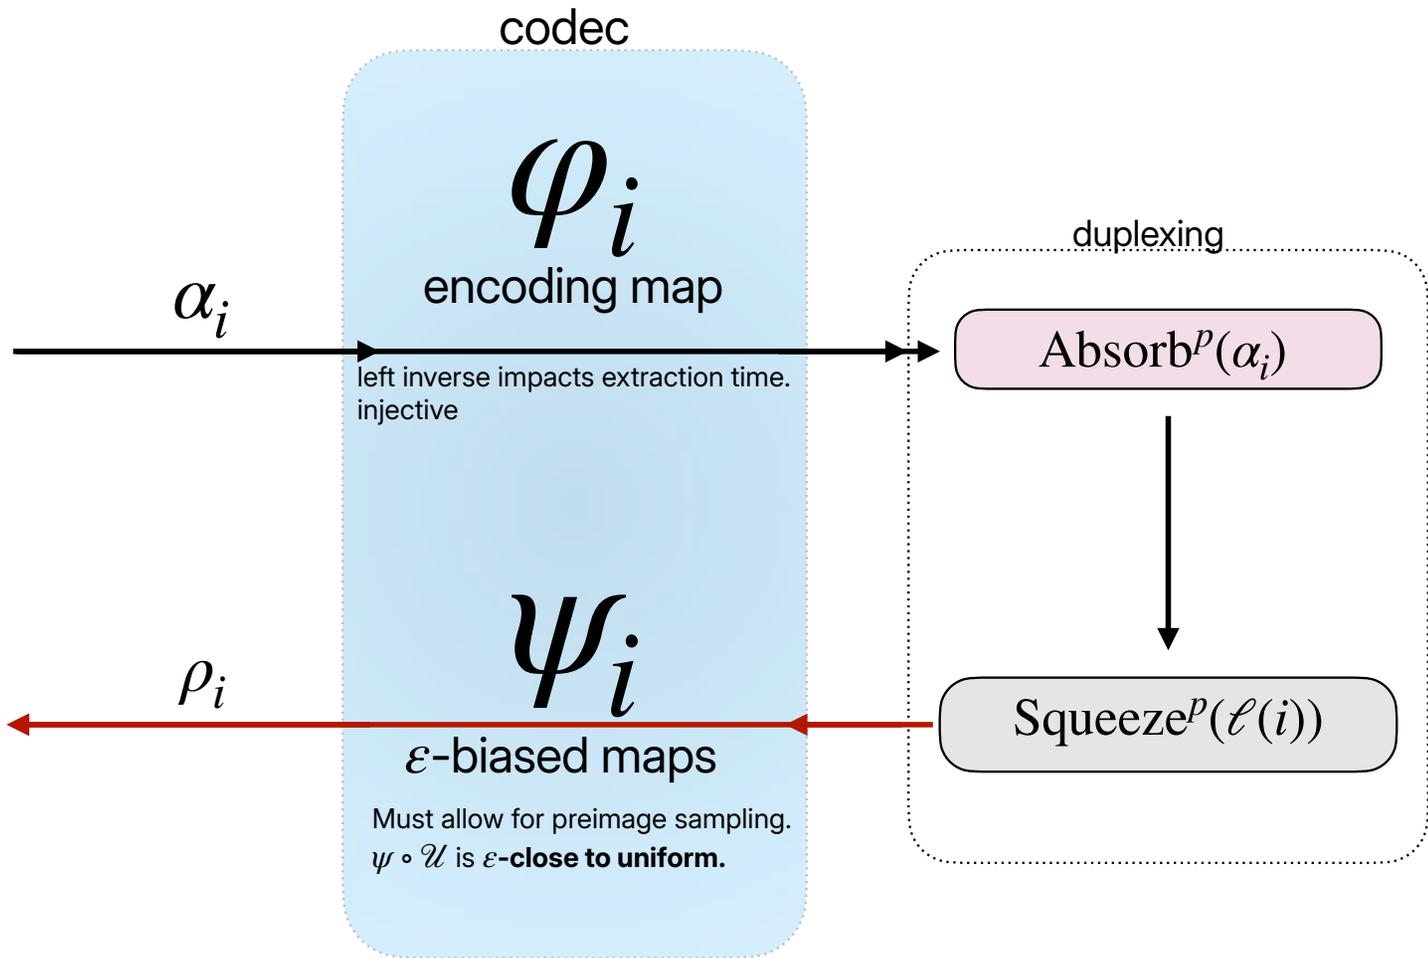
$$b_1 \leftarrow \mathcal{A}^{(p,p^{-1})[\mu]}(\pi)$$

2. We need to translate the positions to be re-programmed

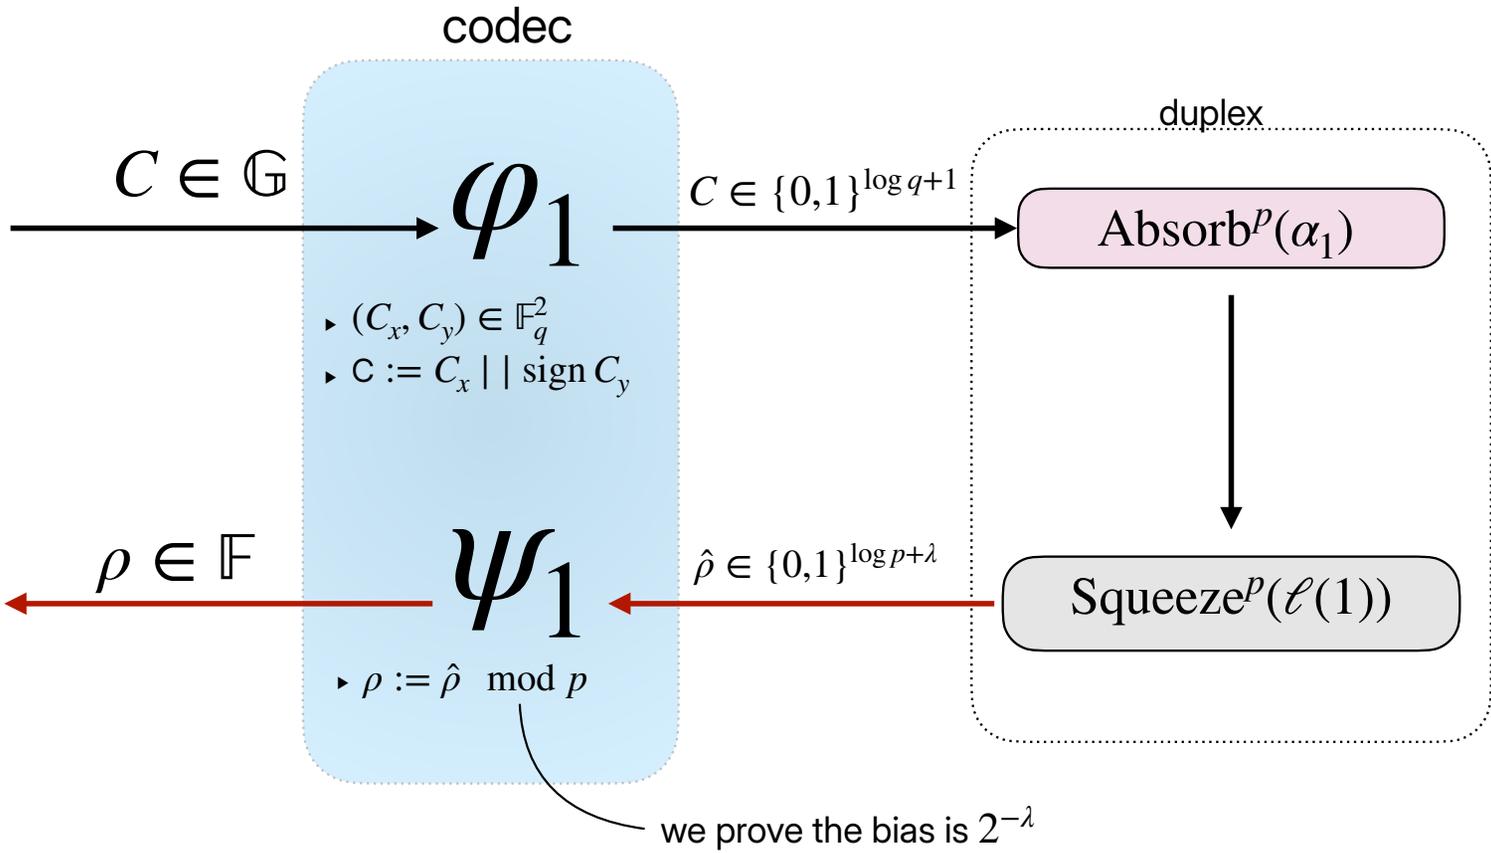
Converting from/to the permutation domain

Codecs

Map from/to permutation domain



Example:
A codec for Schnorr proofs



knowledge soundness:

$$\epsilon_{\text{srsnd}} + O\left(\frac{t^2}{|\Sigma|^c}\right) + O(t) \cdot \max_{i \in [k]} \epsilon_{\text{cdc},i}$$

(depends on the adversarial queries)

Actual security bounds

Taking into account **codecs**, soundness incurs into an additive loss.

zero knowledge:

$$\epsilon_{\text{hvzk}} + O\left(\frac{t}{|\Sigma|^{\min(c,\delta)}}\right) + \sum_{i \in [k]} \epsilon_{\text{cdc},i}$$

(only needed for the programmed queries)

Fiat-Shamir

Oracle $p \Rightarrow$ Real-world p

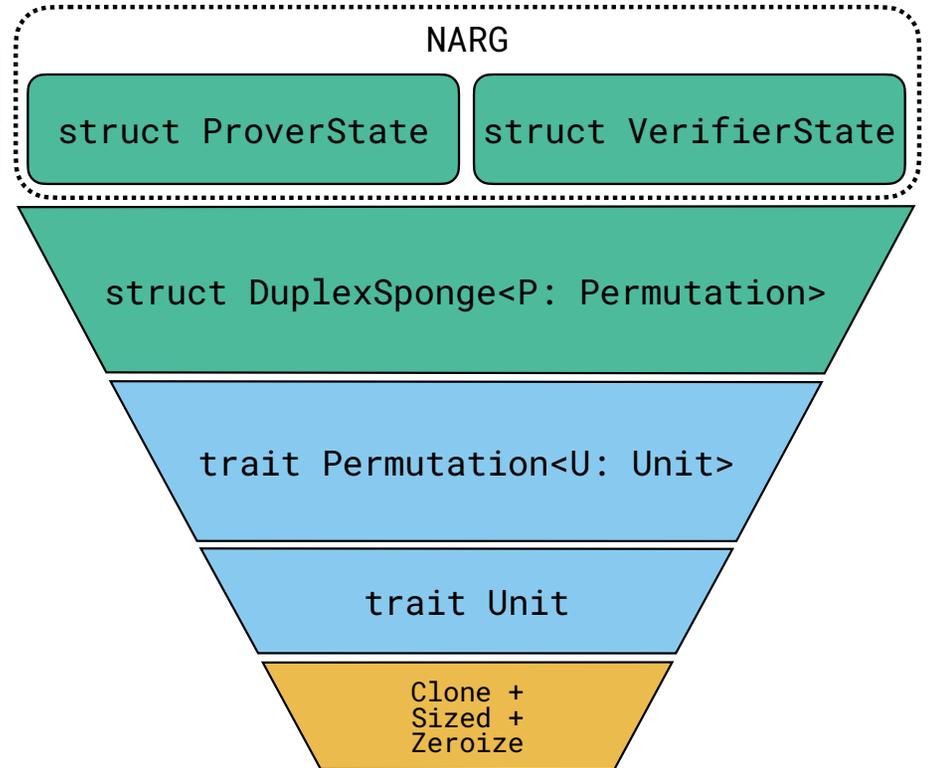
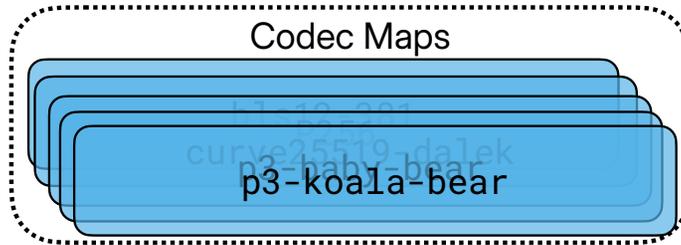
Fiat-Shamir

statistical soundness of IP
+
correlation intractability

Negative results:
[Bar01, GK03, BBHMR19, KRS25]

[IKRR16, CCRR18, HL18, CCHLRRW19, PS19,
LVW19, GJJM19, BFJKS19, LNPT19, CKU20, LV20,
BKM20, JKKZ20, CLMQ20, LNPY20, JJ21, HLR21,
CJJ21a, CJJ21b, LV22, HJKS22, GLS22,
BCHKLPR22, KLV23, CGJJZ23, DJJ24, IL25, ...]

Implementation



Implementation

The software stack of github.com/arkworks-rs/spongefish

Implementation: some perks

```
struct ProverState
```

```
pub fn rng(&mut self) -> &mut (impl CryptoRng + RngCore)
```

Provide the private random coins of the prover using also the witness' entropy.

```
pub fn narg_string(&self) -> &[u8]
```

Offer (inter-operable) proof serialization.

```
!Clone, !Copy
```

Do not implement copy to prevent accidental leaks.

Example: Schnorr proofs

$$\text{DL}(G, X) = \text{NIZK}\{x: xG = X\}$$

$$k \leftarrow \mathbb{F}_p$$
$$K = kG$$

K

```
let k = G::ScalarField::rand(prover_state.rng());  
let K = P * k;
```

```
prover_state.prover_message(&K)?;
```

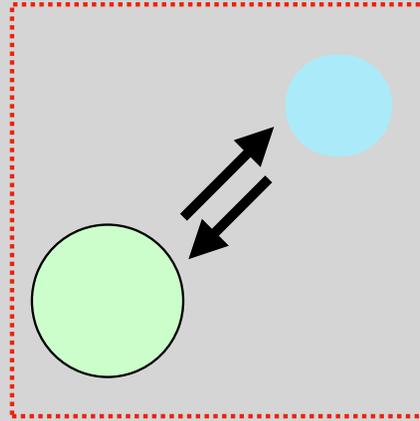
c

```
let c = prover_state.verifier_message::<G::ScalarField>()?;
```

$$r = cx + k$$

```
let r = k + c * x;  
prover_state.prover_message(&r)?;
```

```
Ok(prover_state.narg_string())
```



ia.cr/2025/536



arkworks.rs/spongefish



sigma.zkproof.org/flat-shamir

A Fiat–Shamir transformation from Duplex Sponges

Open problems

Universal composability

Is it possible to prove universal composability in the global "ideal permutation" model (in analogy to GROM)?

Diagonalization attacks

How does this interact with recent attacks and fixes for Fiat-Shamir, and what about the duplex sponge transformation?