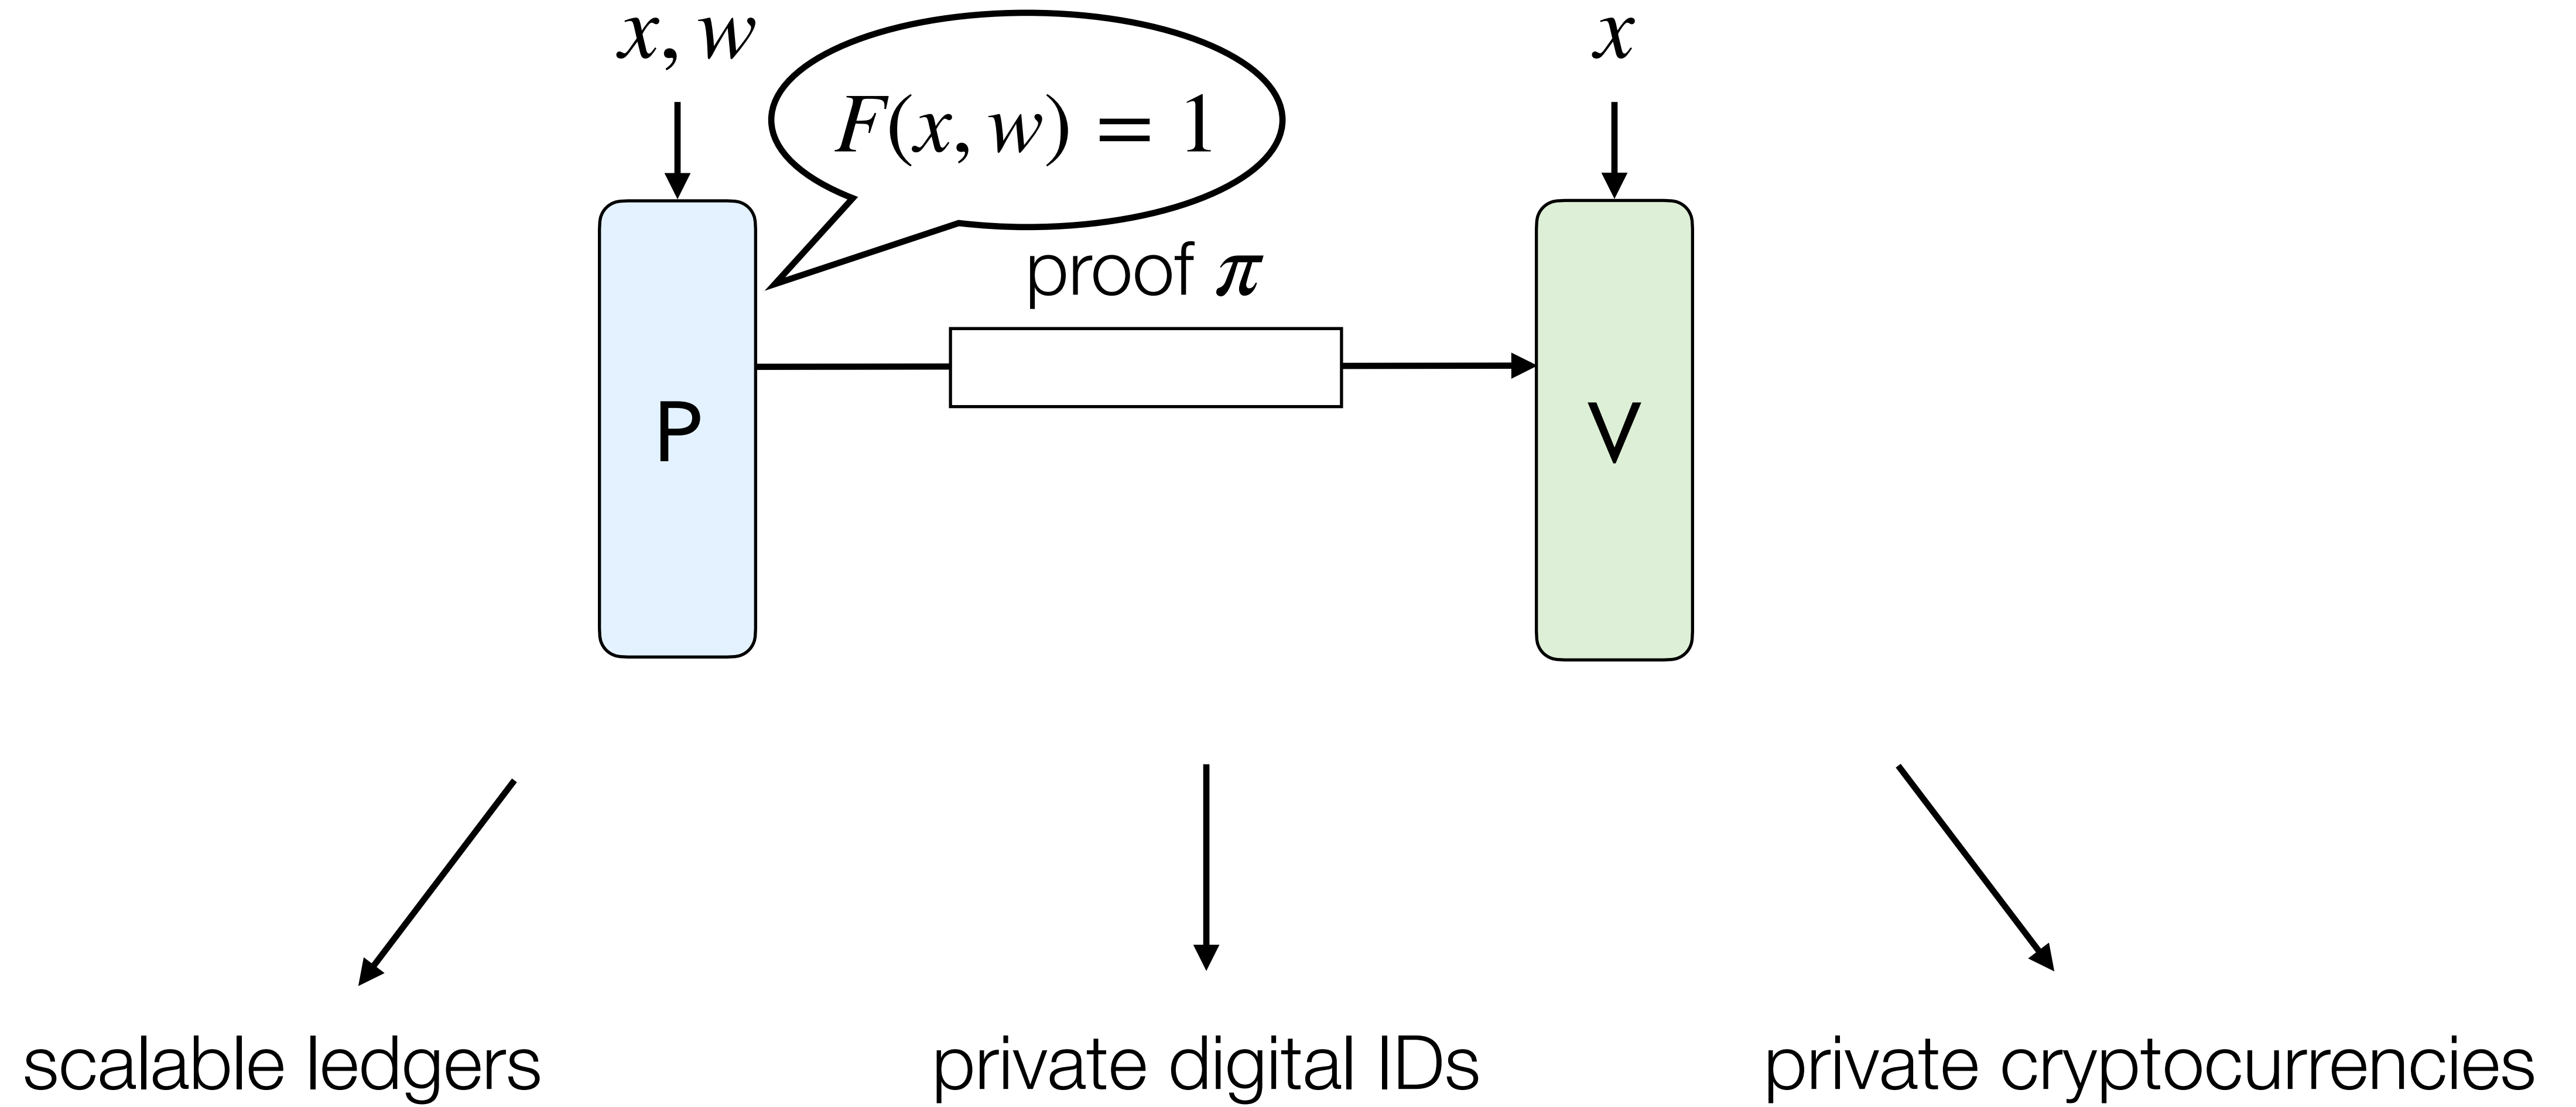


The Quest for Optimal IOPs

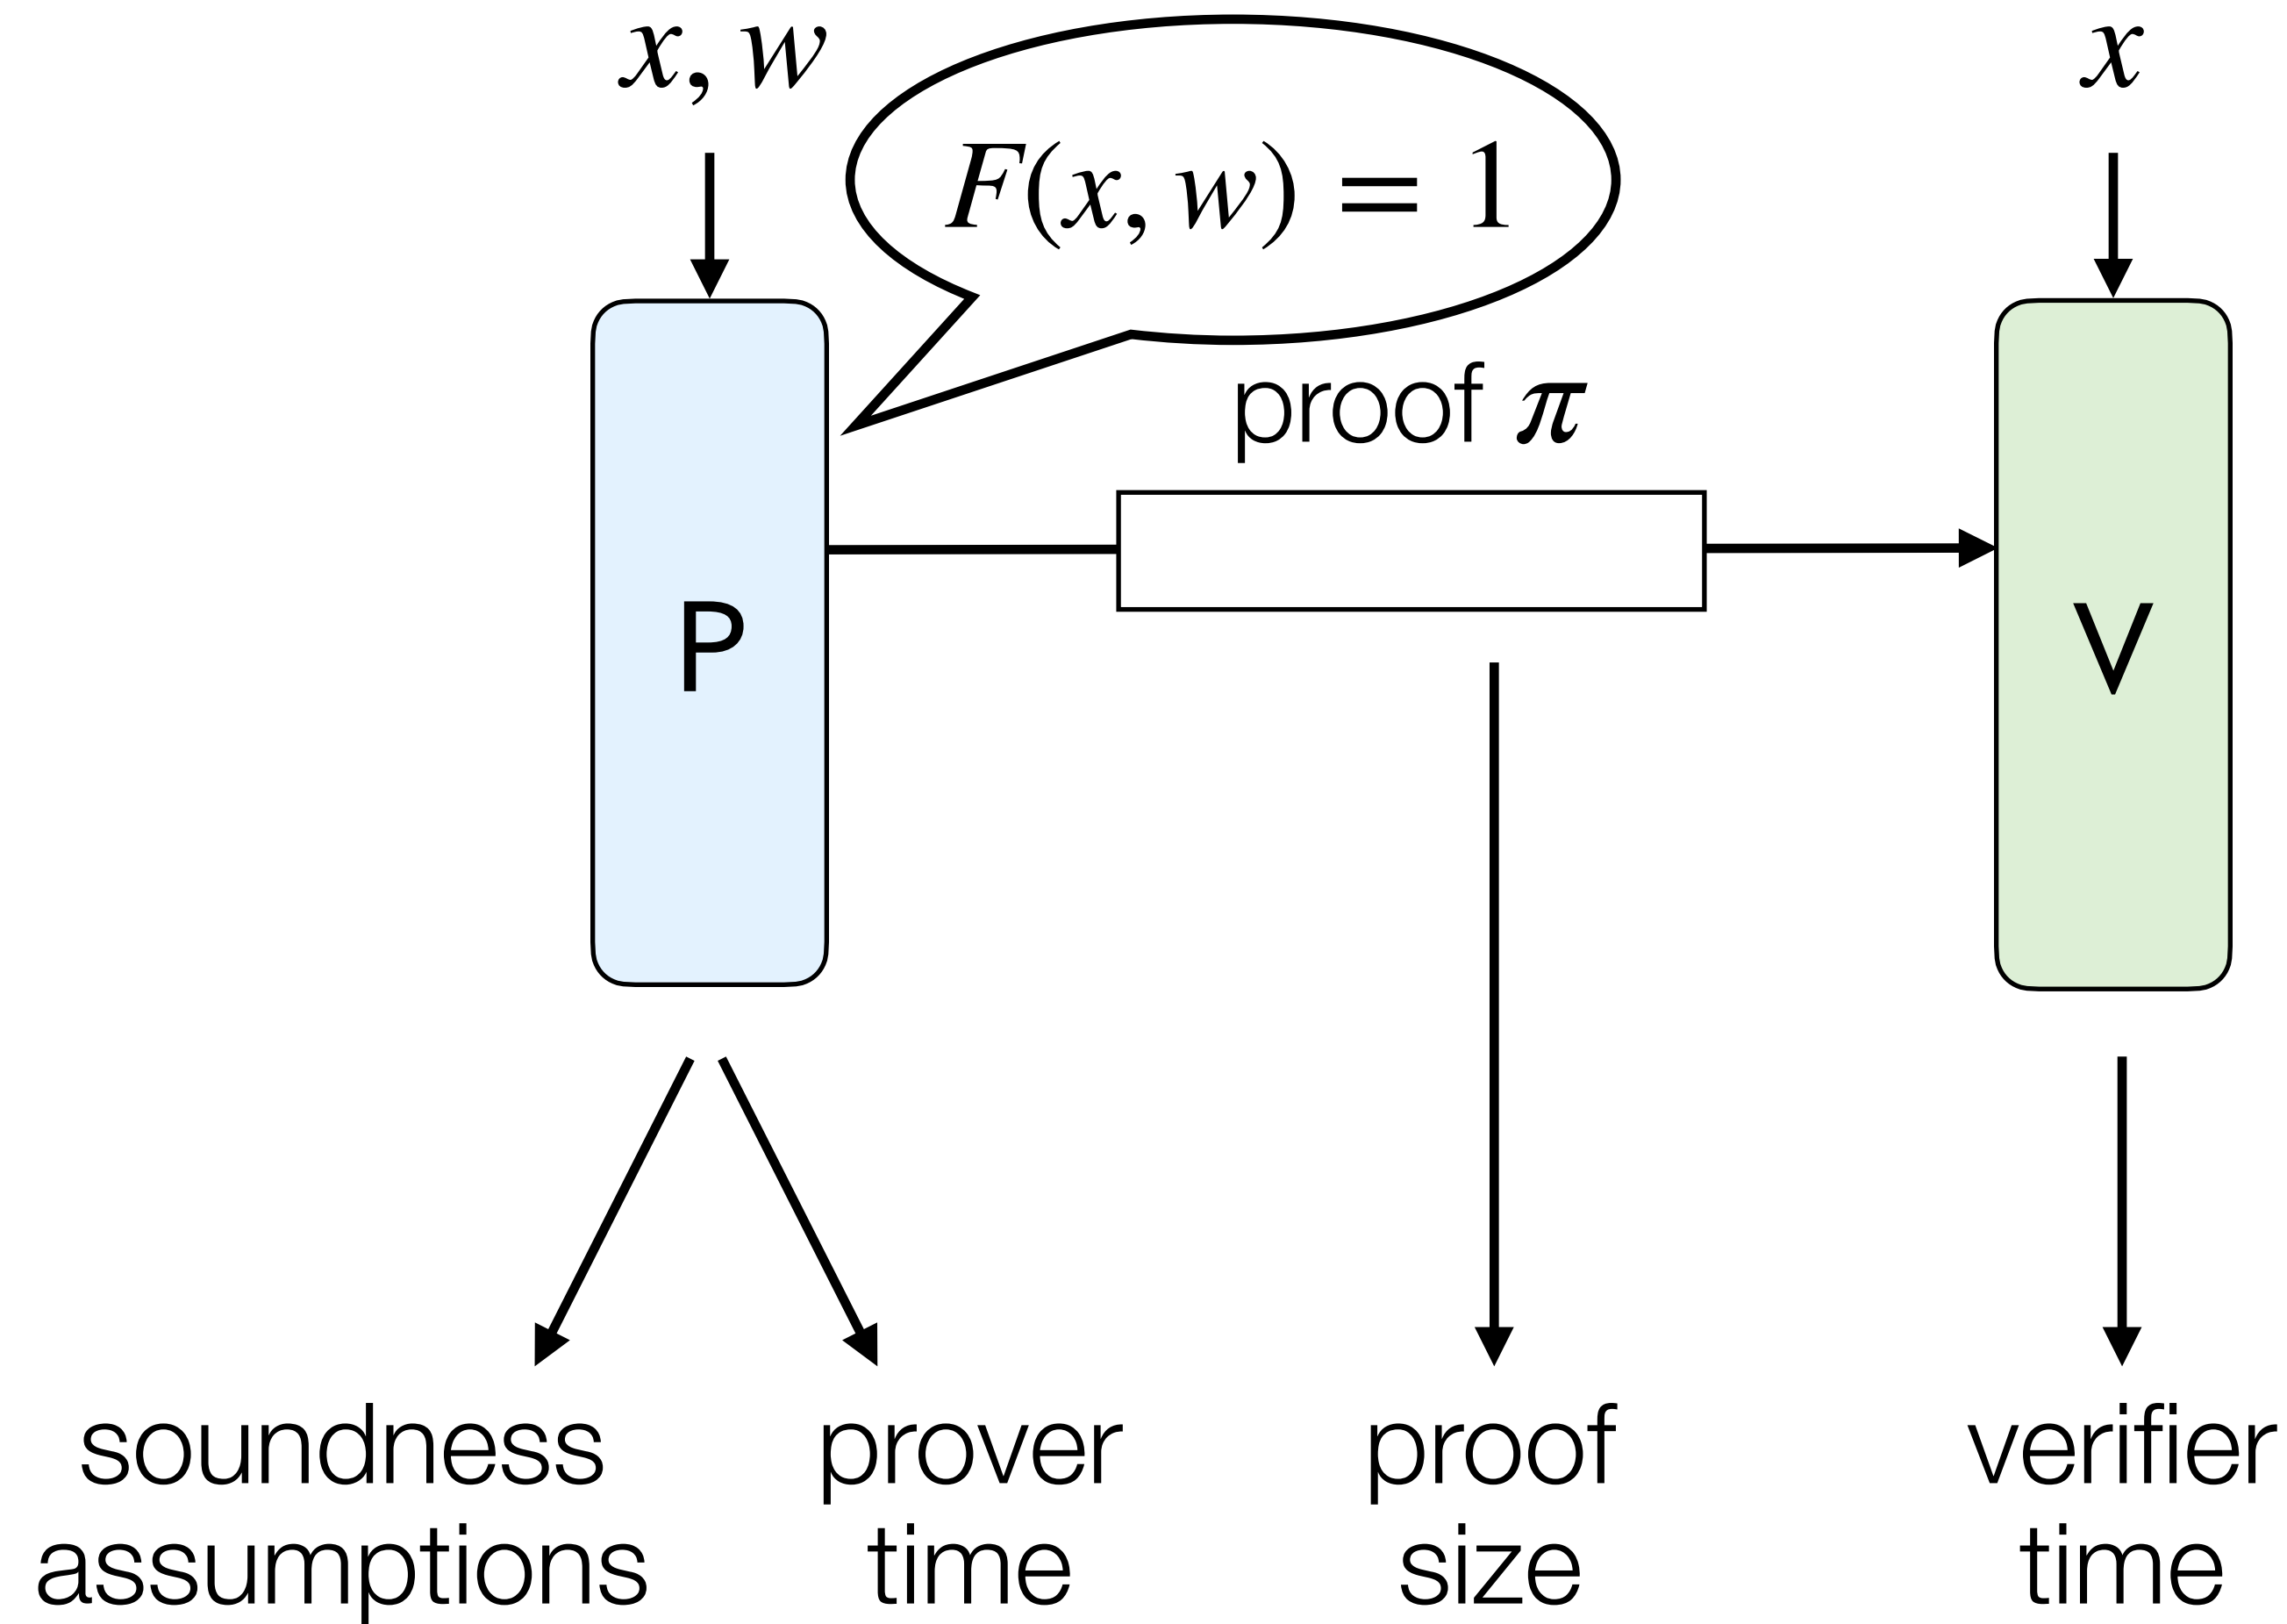
Pratyush Mishra

University of Pennsylvania

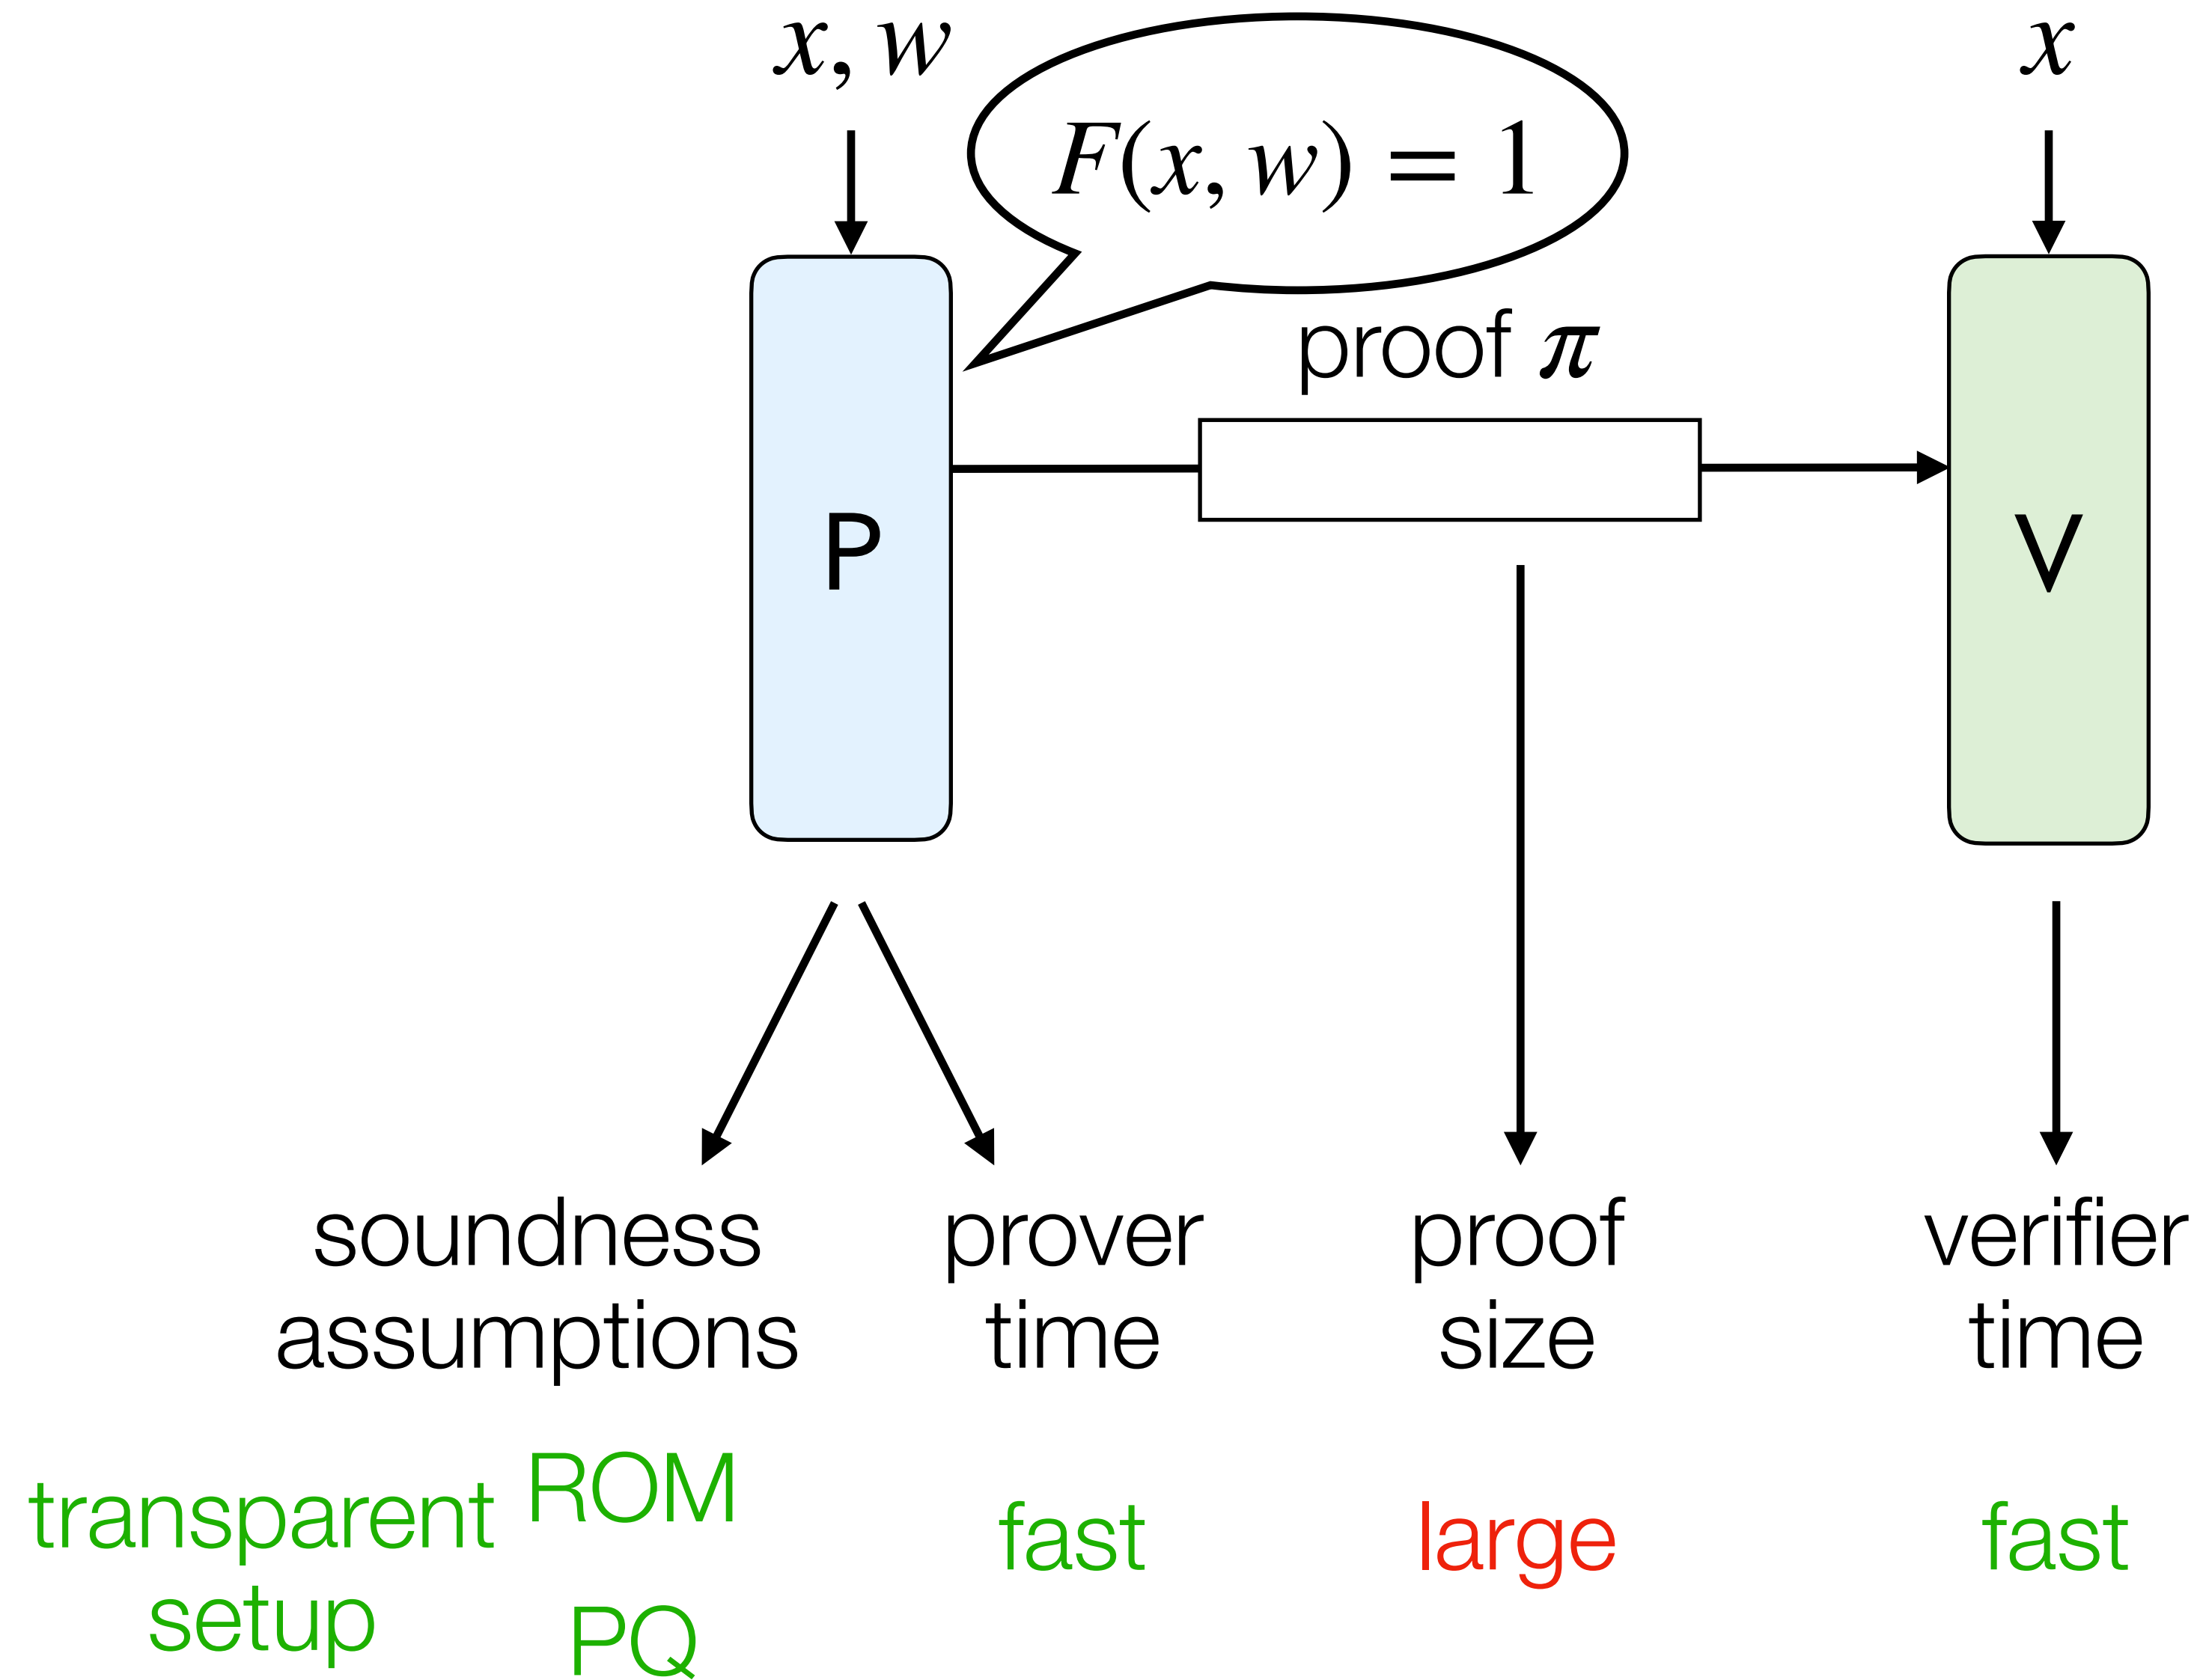
SNARKs




What Properties Do We Care About?




SNARKs from IOPs [K92, Mic94, BCS16]



Interactive Oracle Proofs (IOPs) [BCS16, RRR16]



IOPs \rightarrow SNARKs [Mic94, BCS16]



prover
time

proof
size

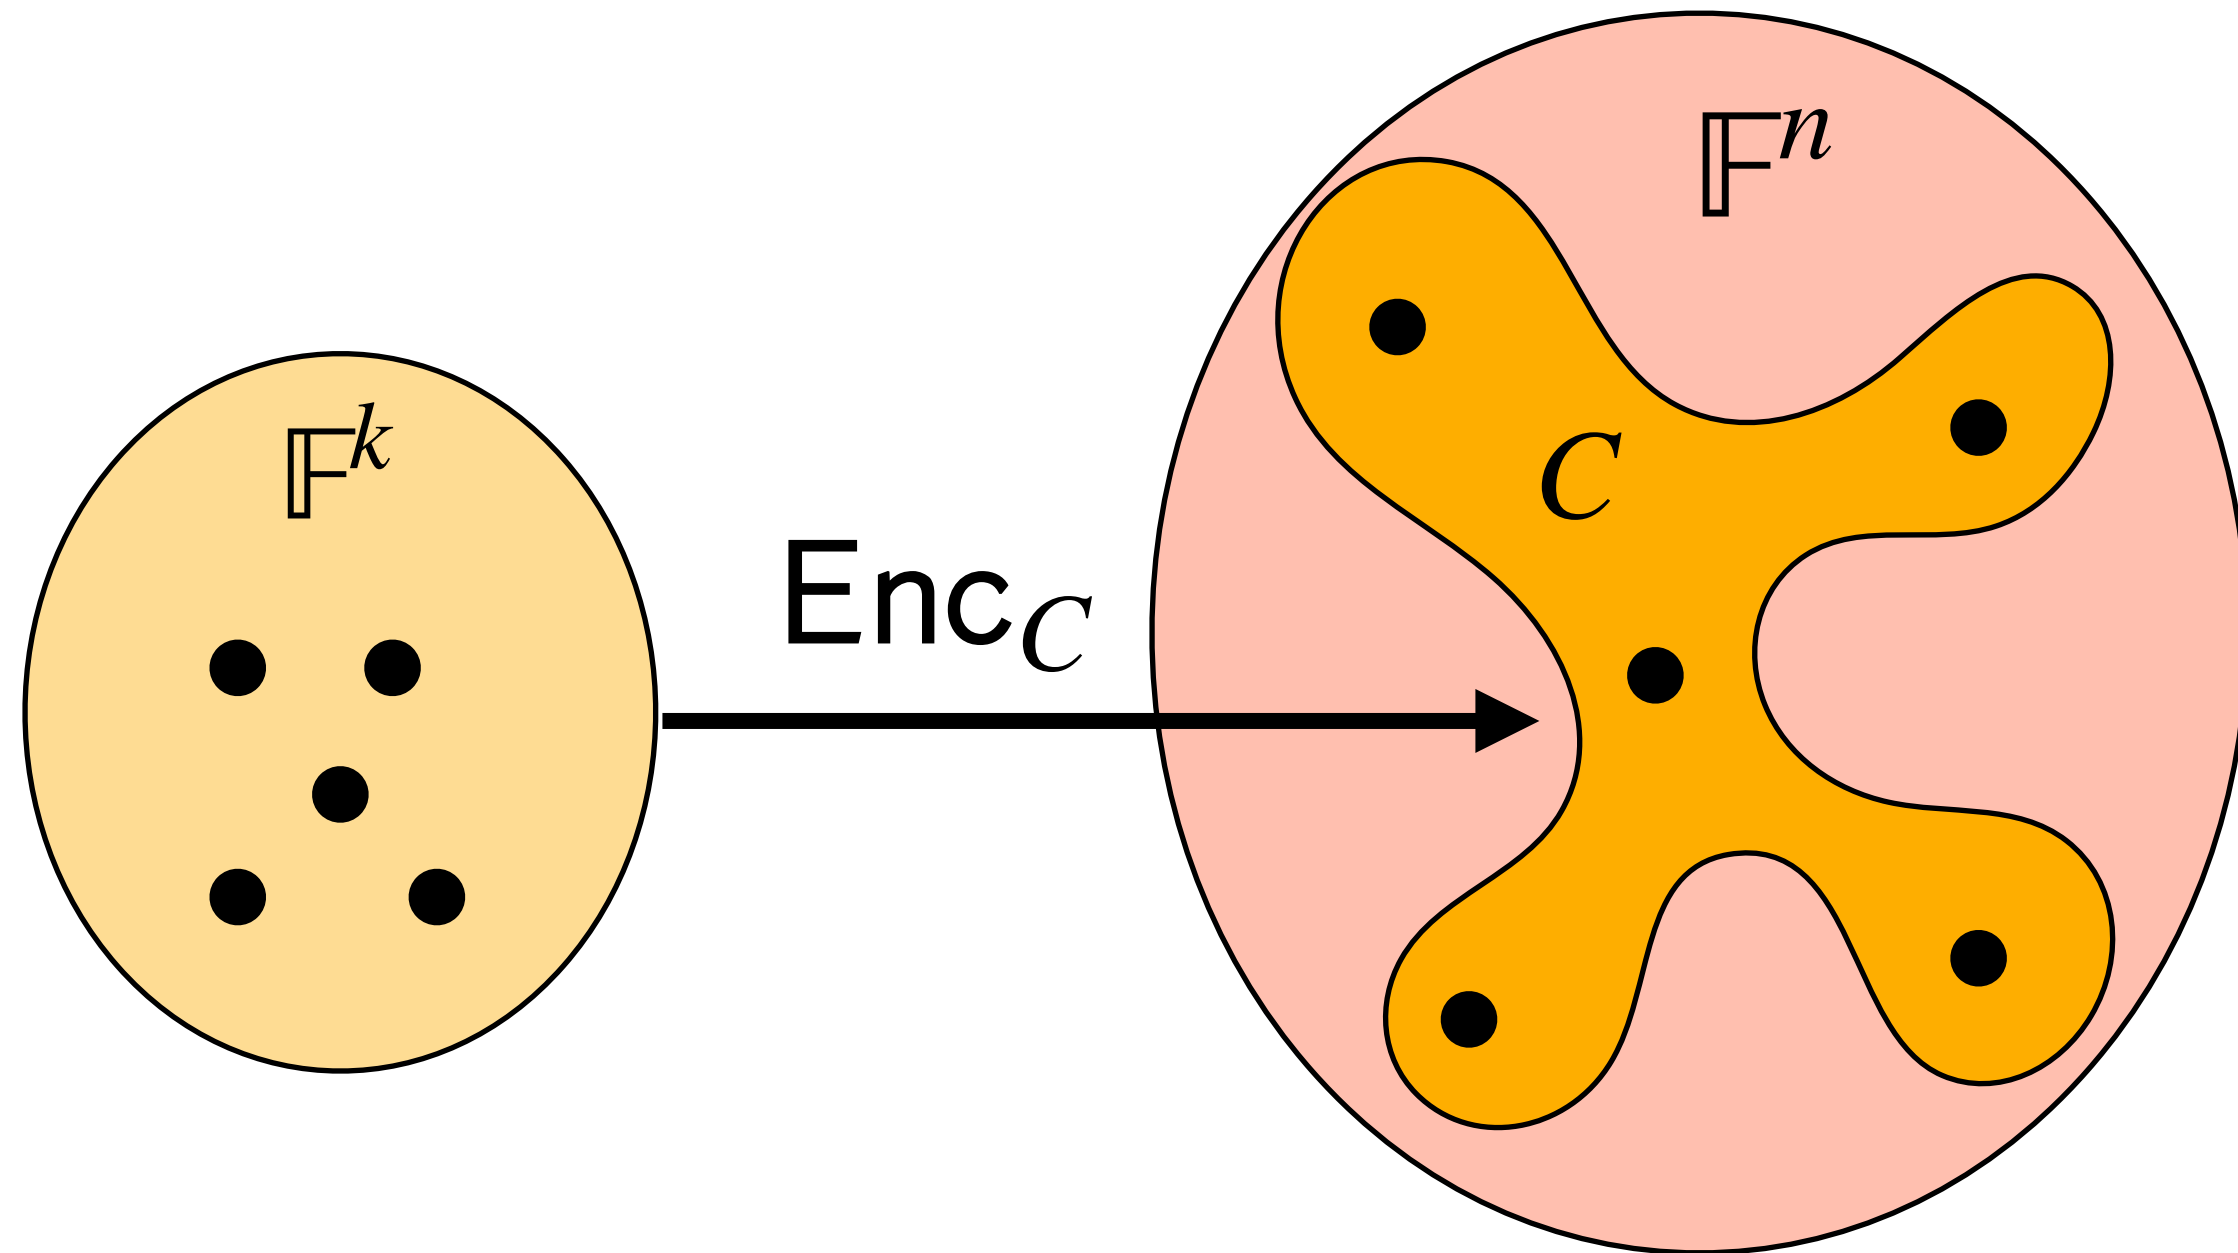
verifier
time

$|P|$

$Q \log n$

$Q \log n$ hashes + $|V|$

Constructing IOPs from Codes



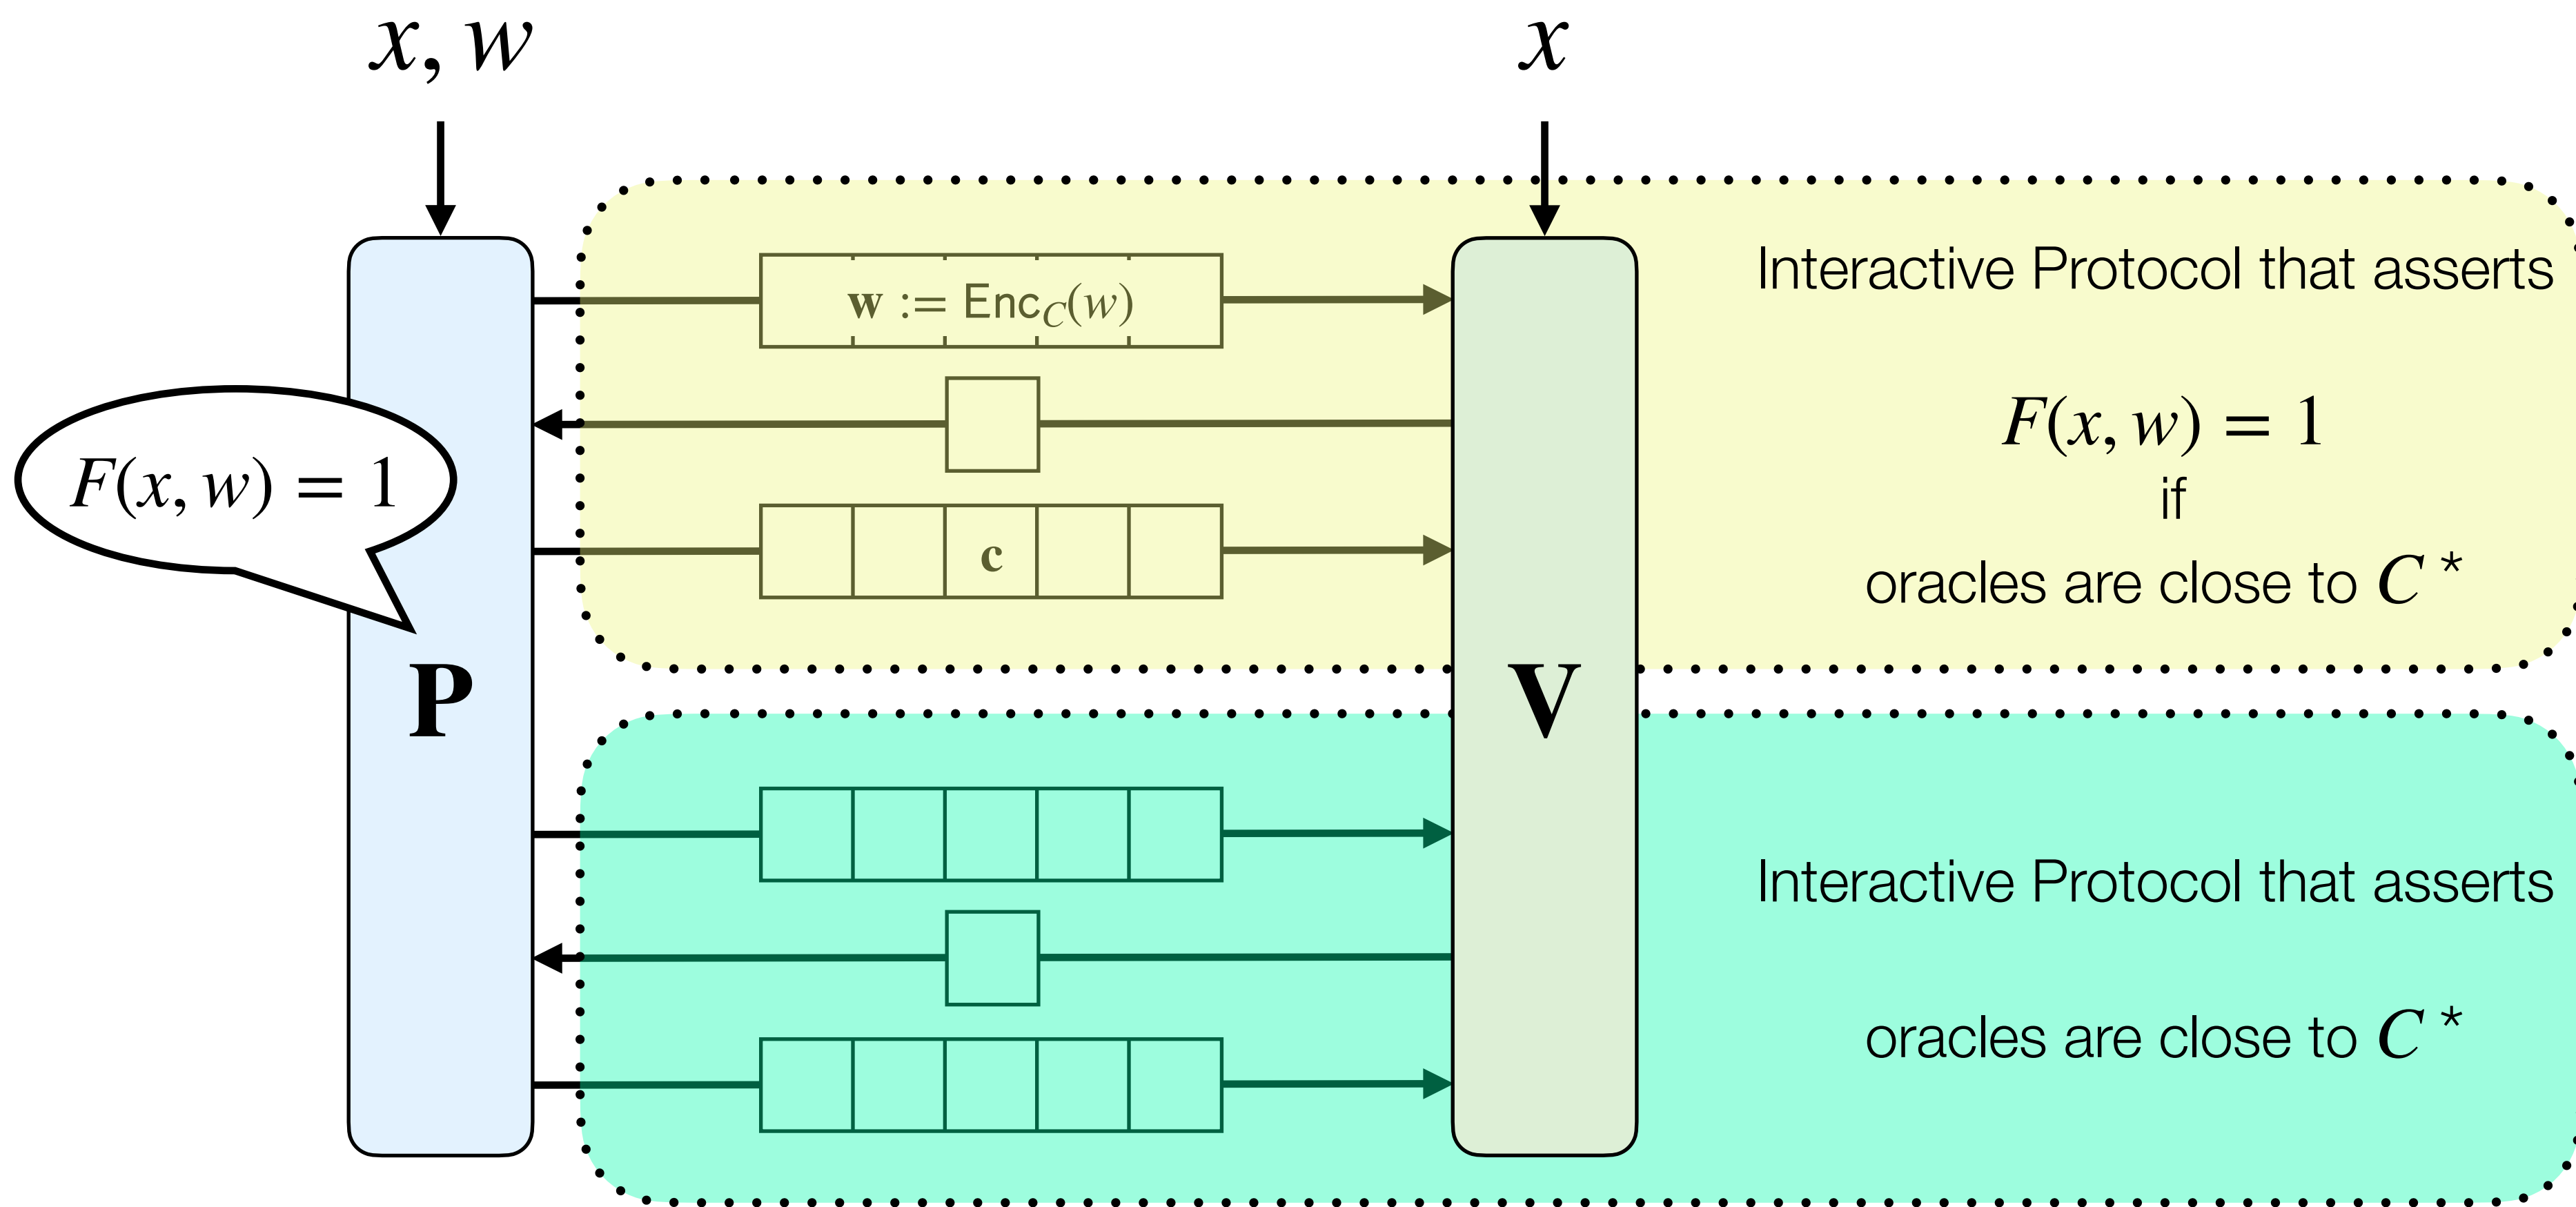
Distance δ : how far apart are points in C ?

Rate $\rho = k/n$: how redundant is C ?

Encoding time T_C : how long does Enc_C take?

Constructing IOPs from Codes

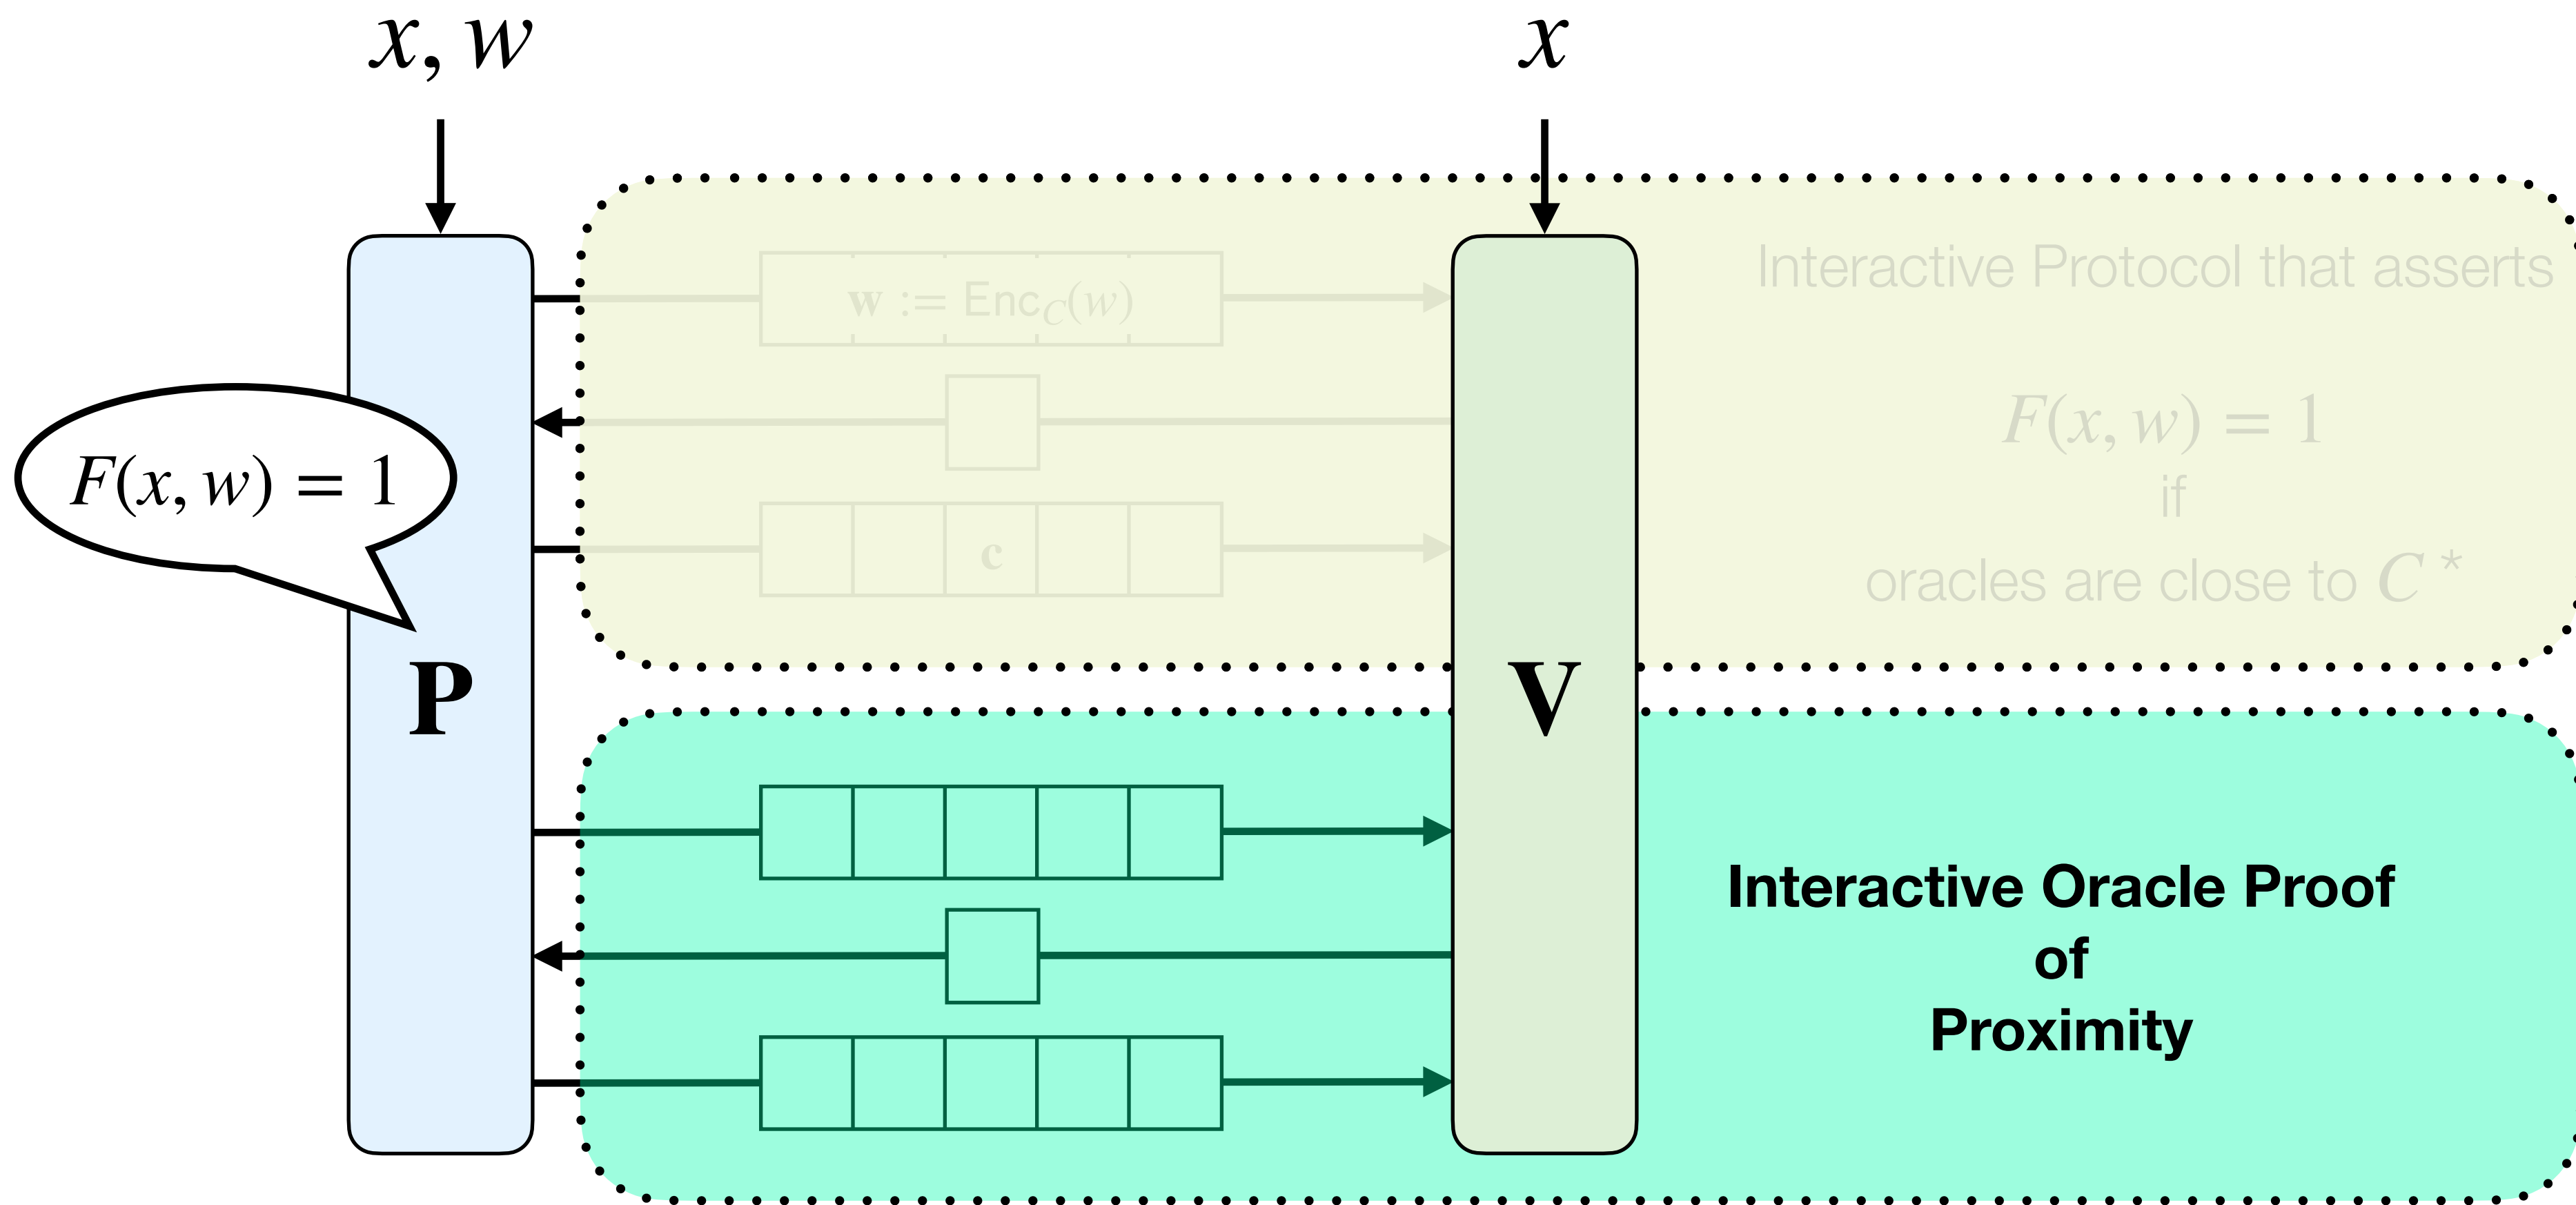
Code C with
 Distance δ
 Rate $\rho = k/n$
 Encoding time T_C



* and (simple) claims about underlying messages are true

Constructing IOPs from Codes

Code C with
 Distance δ
 Rate $\rho = k/n$
 Encoding time T_C



Near-optimal constructions
 [Set20, CBBZ23]
Prover time: $O(T_F + T_C)$
Query complexity: 0 (!)

Prover time: ???
Query complexity: ???

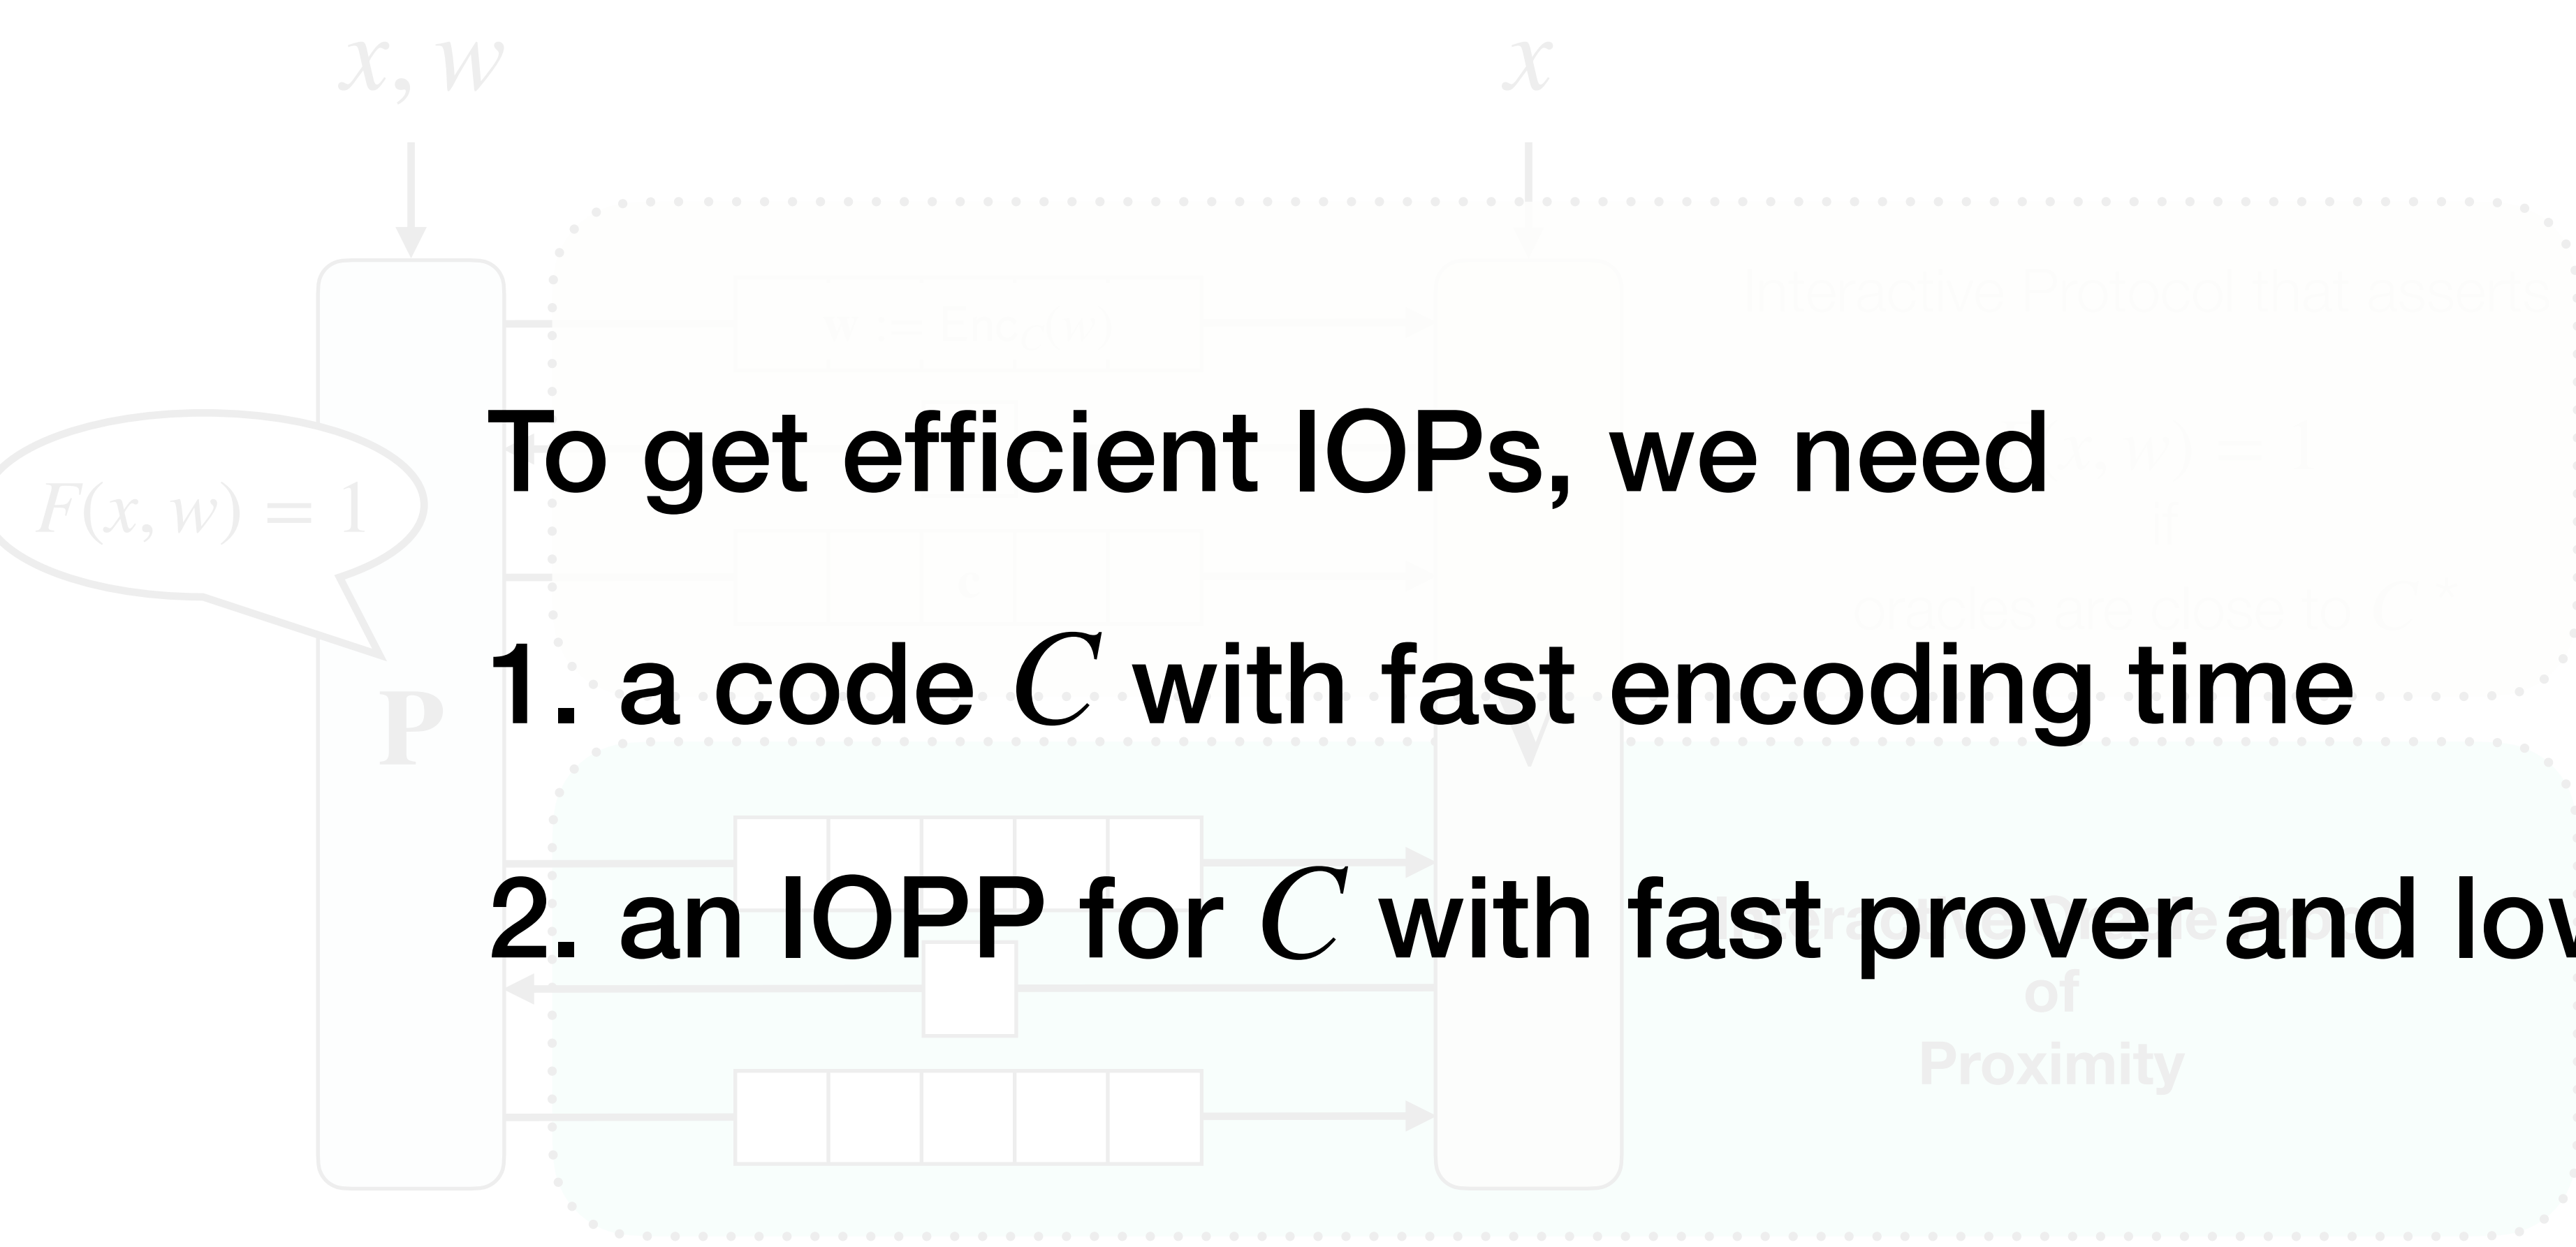
IOP prover time = $O(T_F + T_C + T_{\text{IOPP}})$

IOP query complexity = Q_{IOPP}

* and (simple) claims about underlying messages are true

Constructing IOPs from Codes

Key tool: Code C with
 Distance δ
 Rate $\rho = k/n$
 Encoding time T_C



To get efficient IOPs, we need

1. a code C with fast encoding time

2. an IOPP for C with fast prover and low # queries

Near-optimal constructions
 [Set20, CBBZ23]
 Prover time: $O(T_F + T_C)$
 Query complexity: 0 (!)

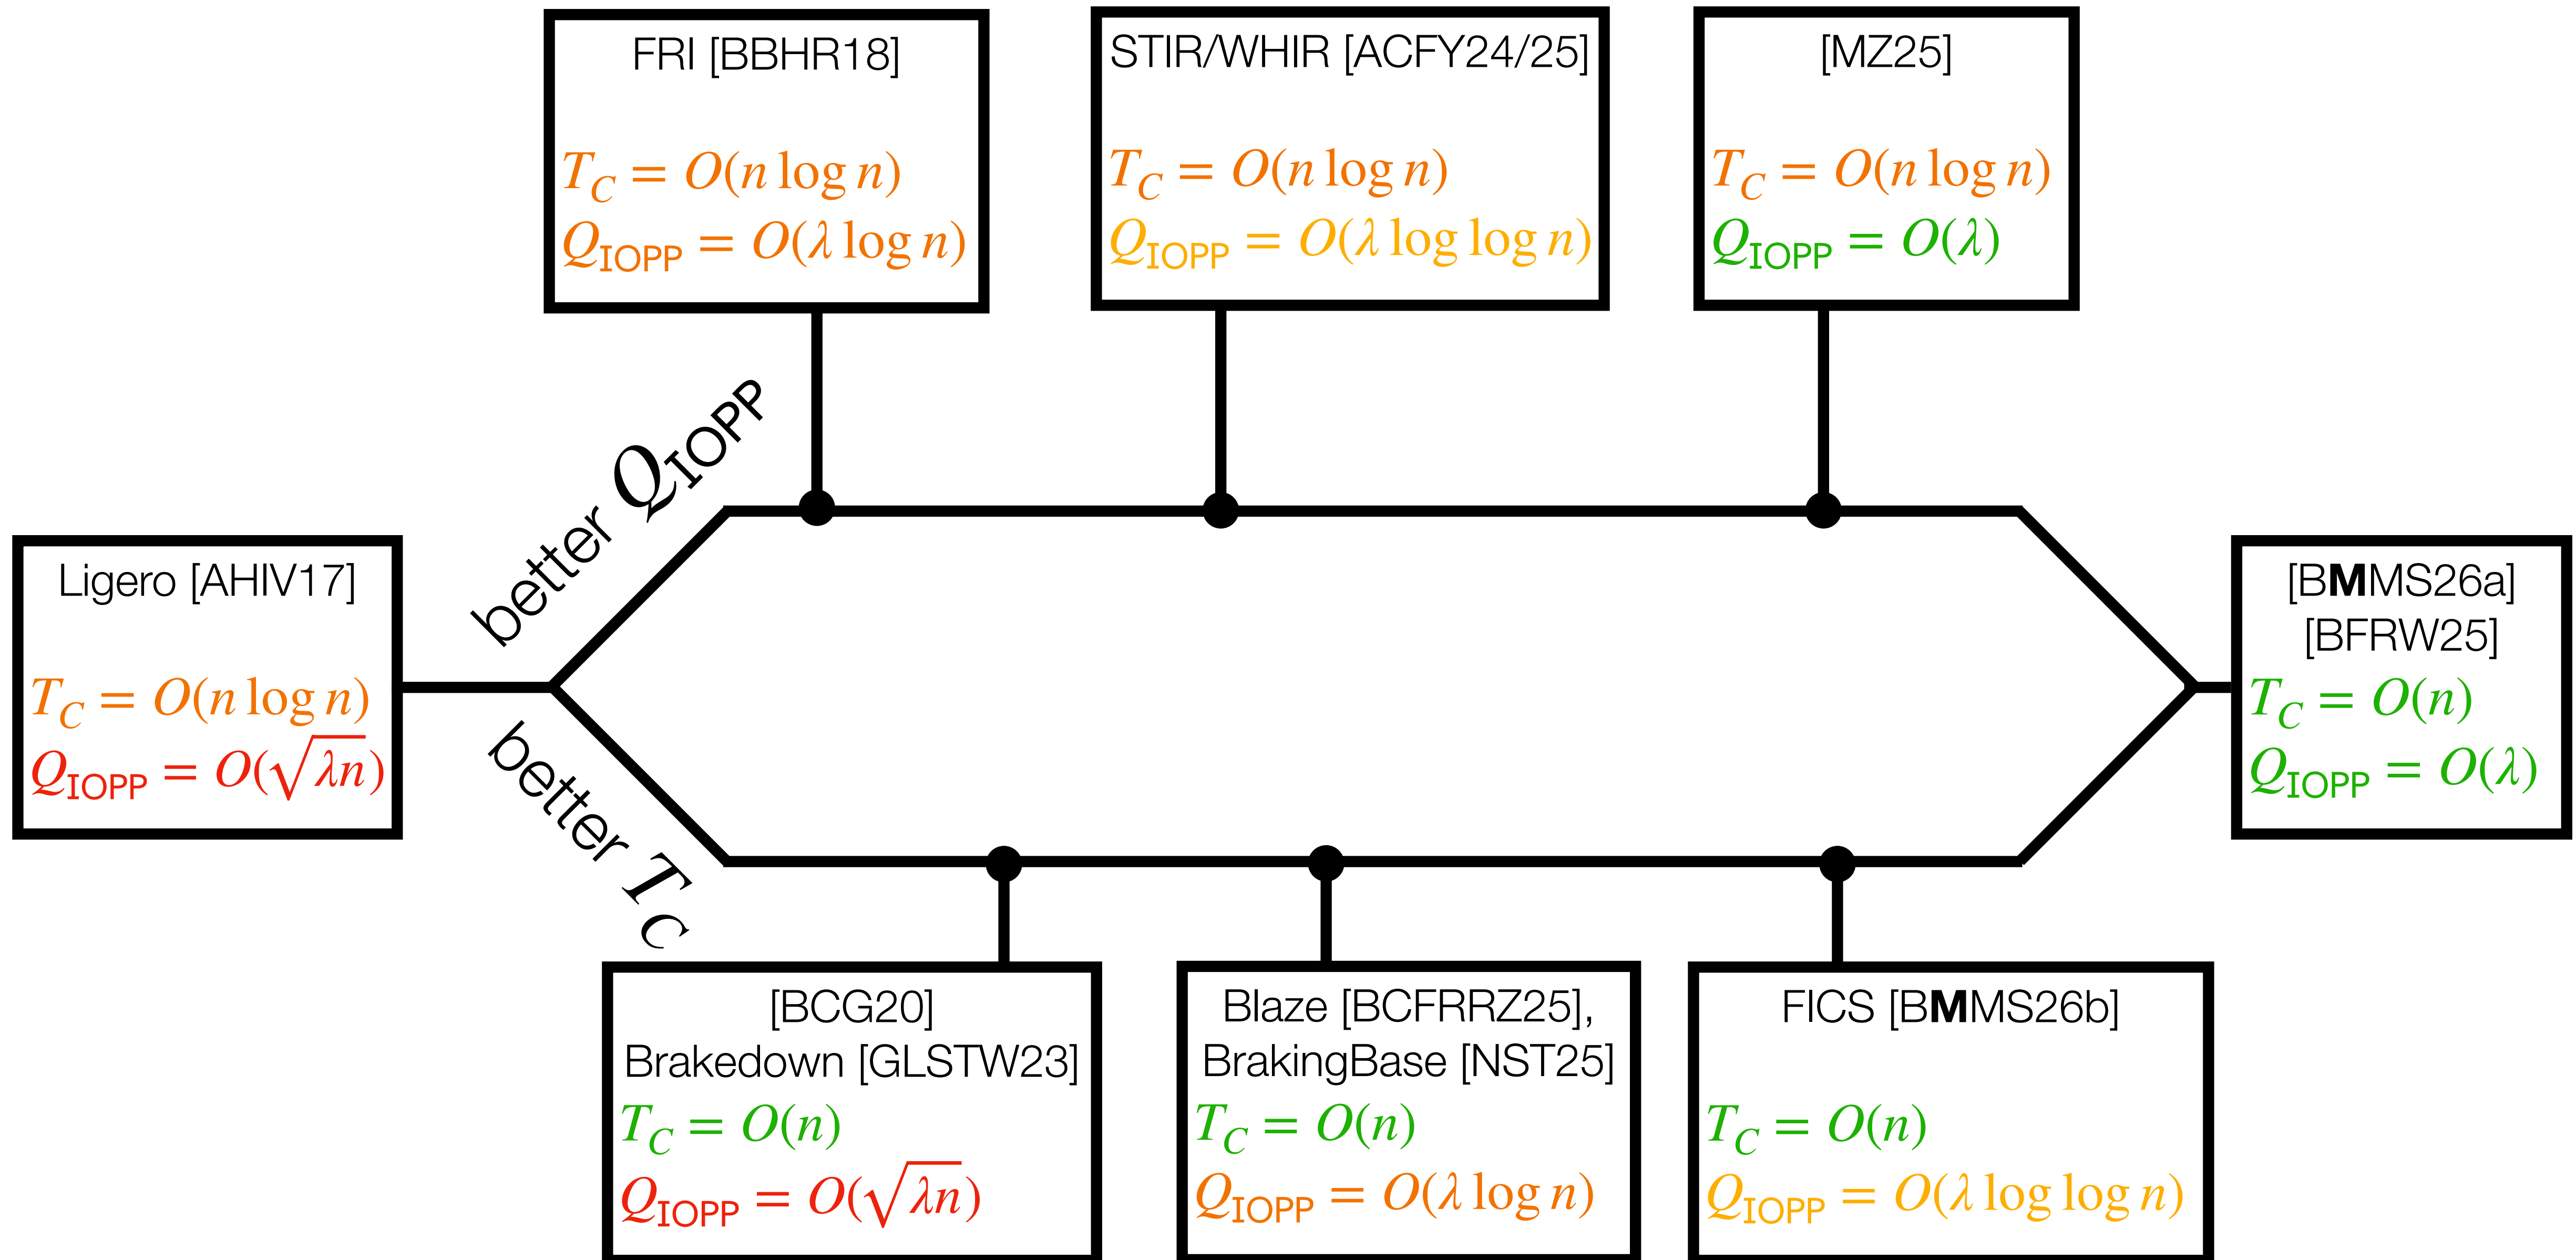
Prover time: ???
 Query complexity: ???

Prover time = $O(T_F + T_C + T_{\text{IOPP}})$

Query complexity = Q_{IOPP}

* and (simple) claims about underlying messages are true

Efficient Codes and IOPPs for them



Chapter 1

IOPPs by Interleaving Codes

Goal:

IOPP with $O(\sqrt{\lambda k})$ queries

Interleaved Codes

Let $C \subseteq \mathbb{F}^r$ be a code with message space \mathbb{F}^s .

The t -wise **interleaving** of C is denoted by C^t

$\mathbf{m}_1 \in \mathbb{F}^s$
$\mathbf{m}_2 \in \mathbb{F}^s$
$\mathbf{m}_3 \in \mathbb{F}^s$
\vdots
$\mathbf{m}_t \in \mathbb{F}^s$

$\xrightarrow{\text{Enc}_{C^t}}$

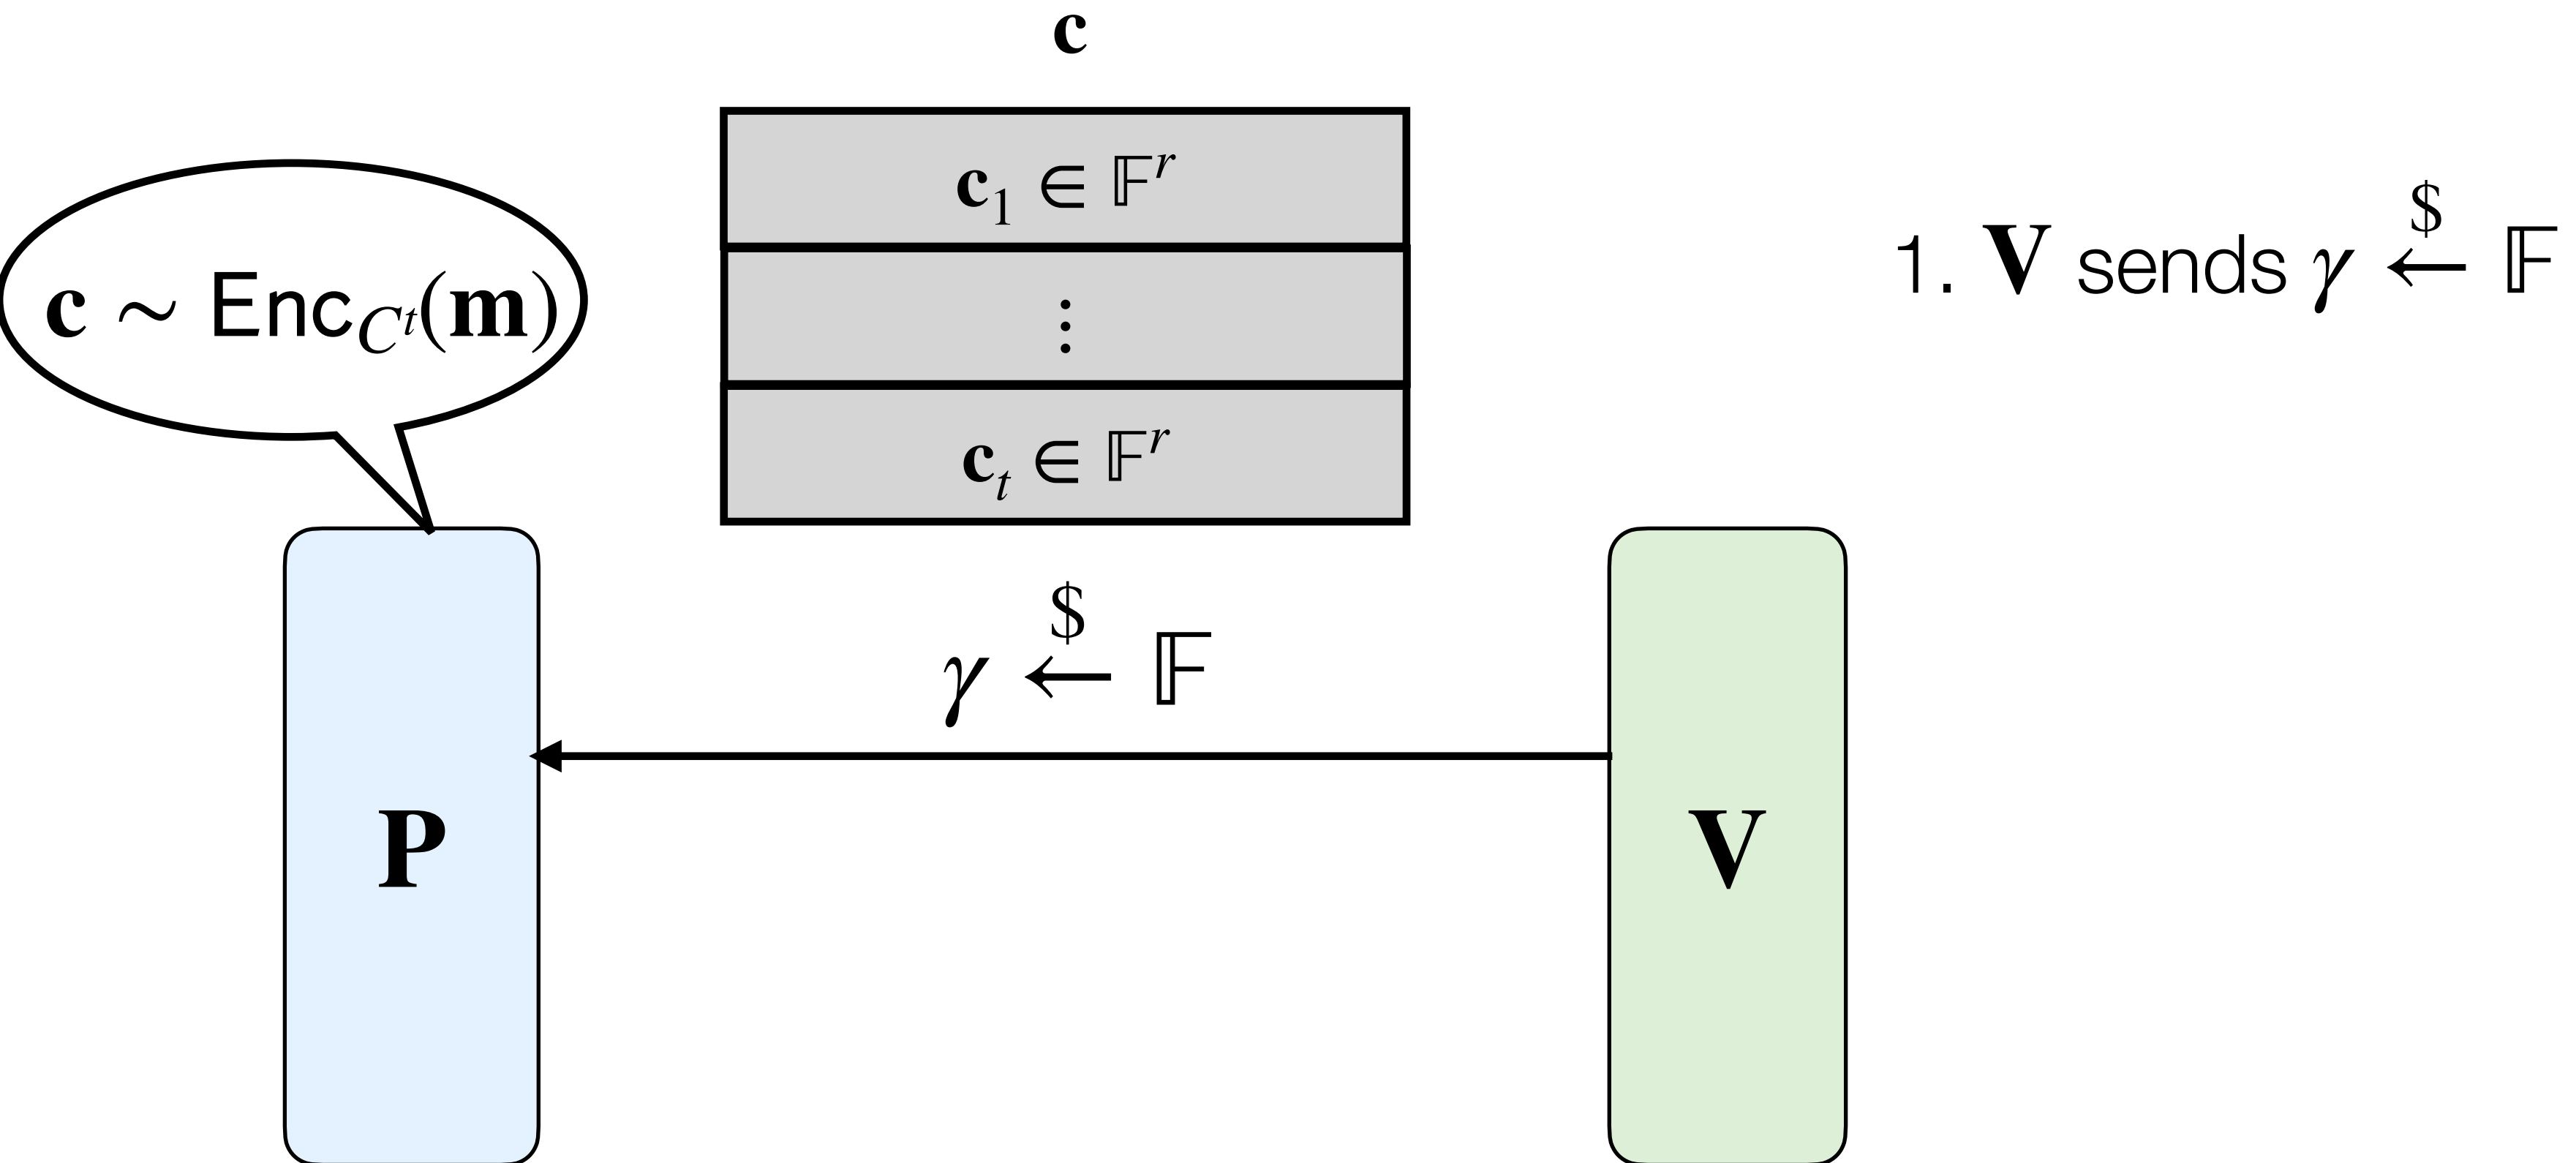
$\mathbf{c}_1 \in \mathbb{F}^r$
$\mathbf{c}_2 \in \mathbb{F}^r$
$\mathbf{c}_3 \in \mathbb{F}^r$
\vdots
$\mathbf{c}_t \in \mathbb{F}^r$

Rate $\rho_{C^t} = s/r = \rho_C$

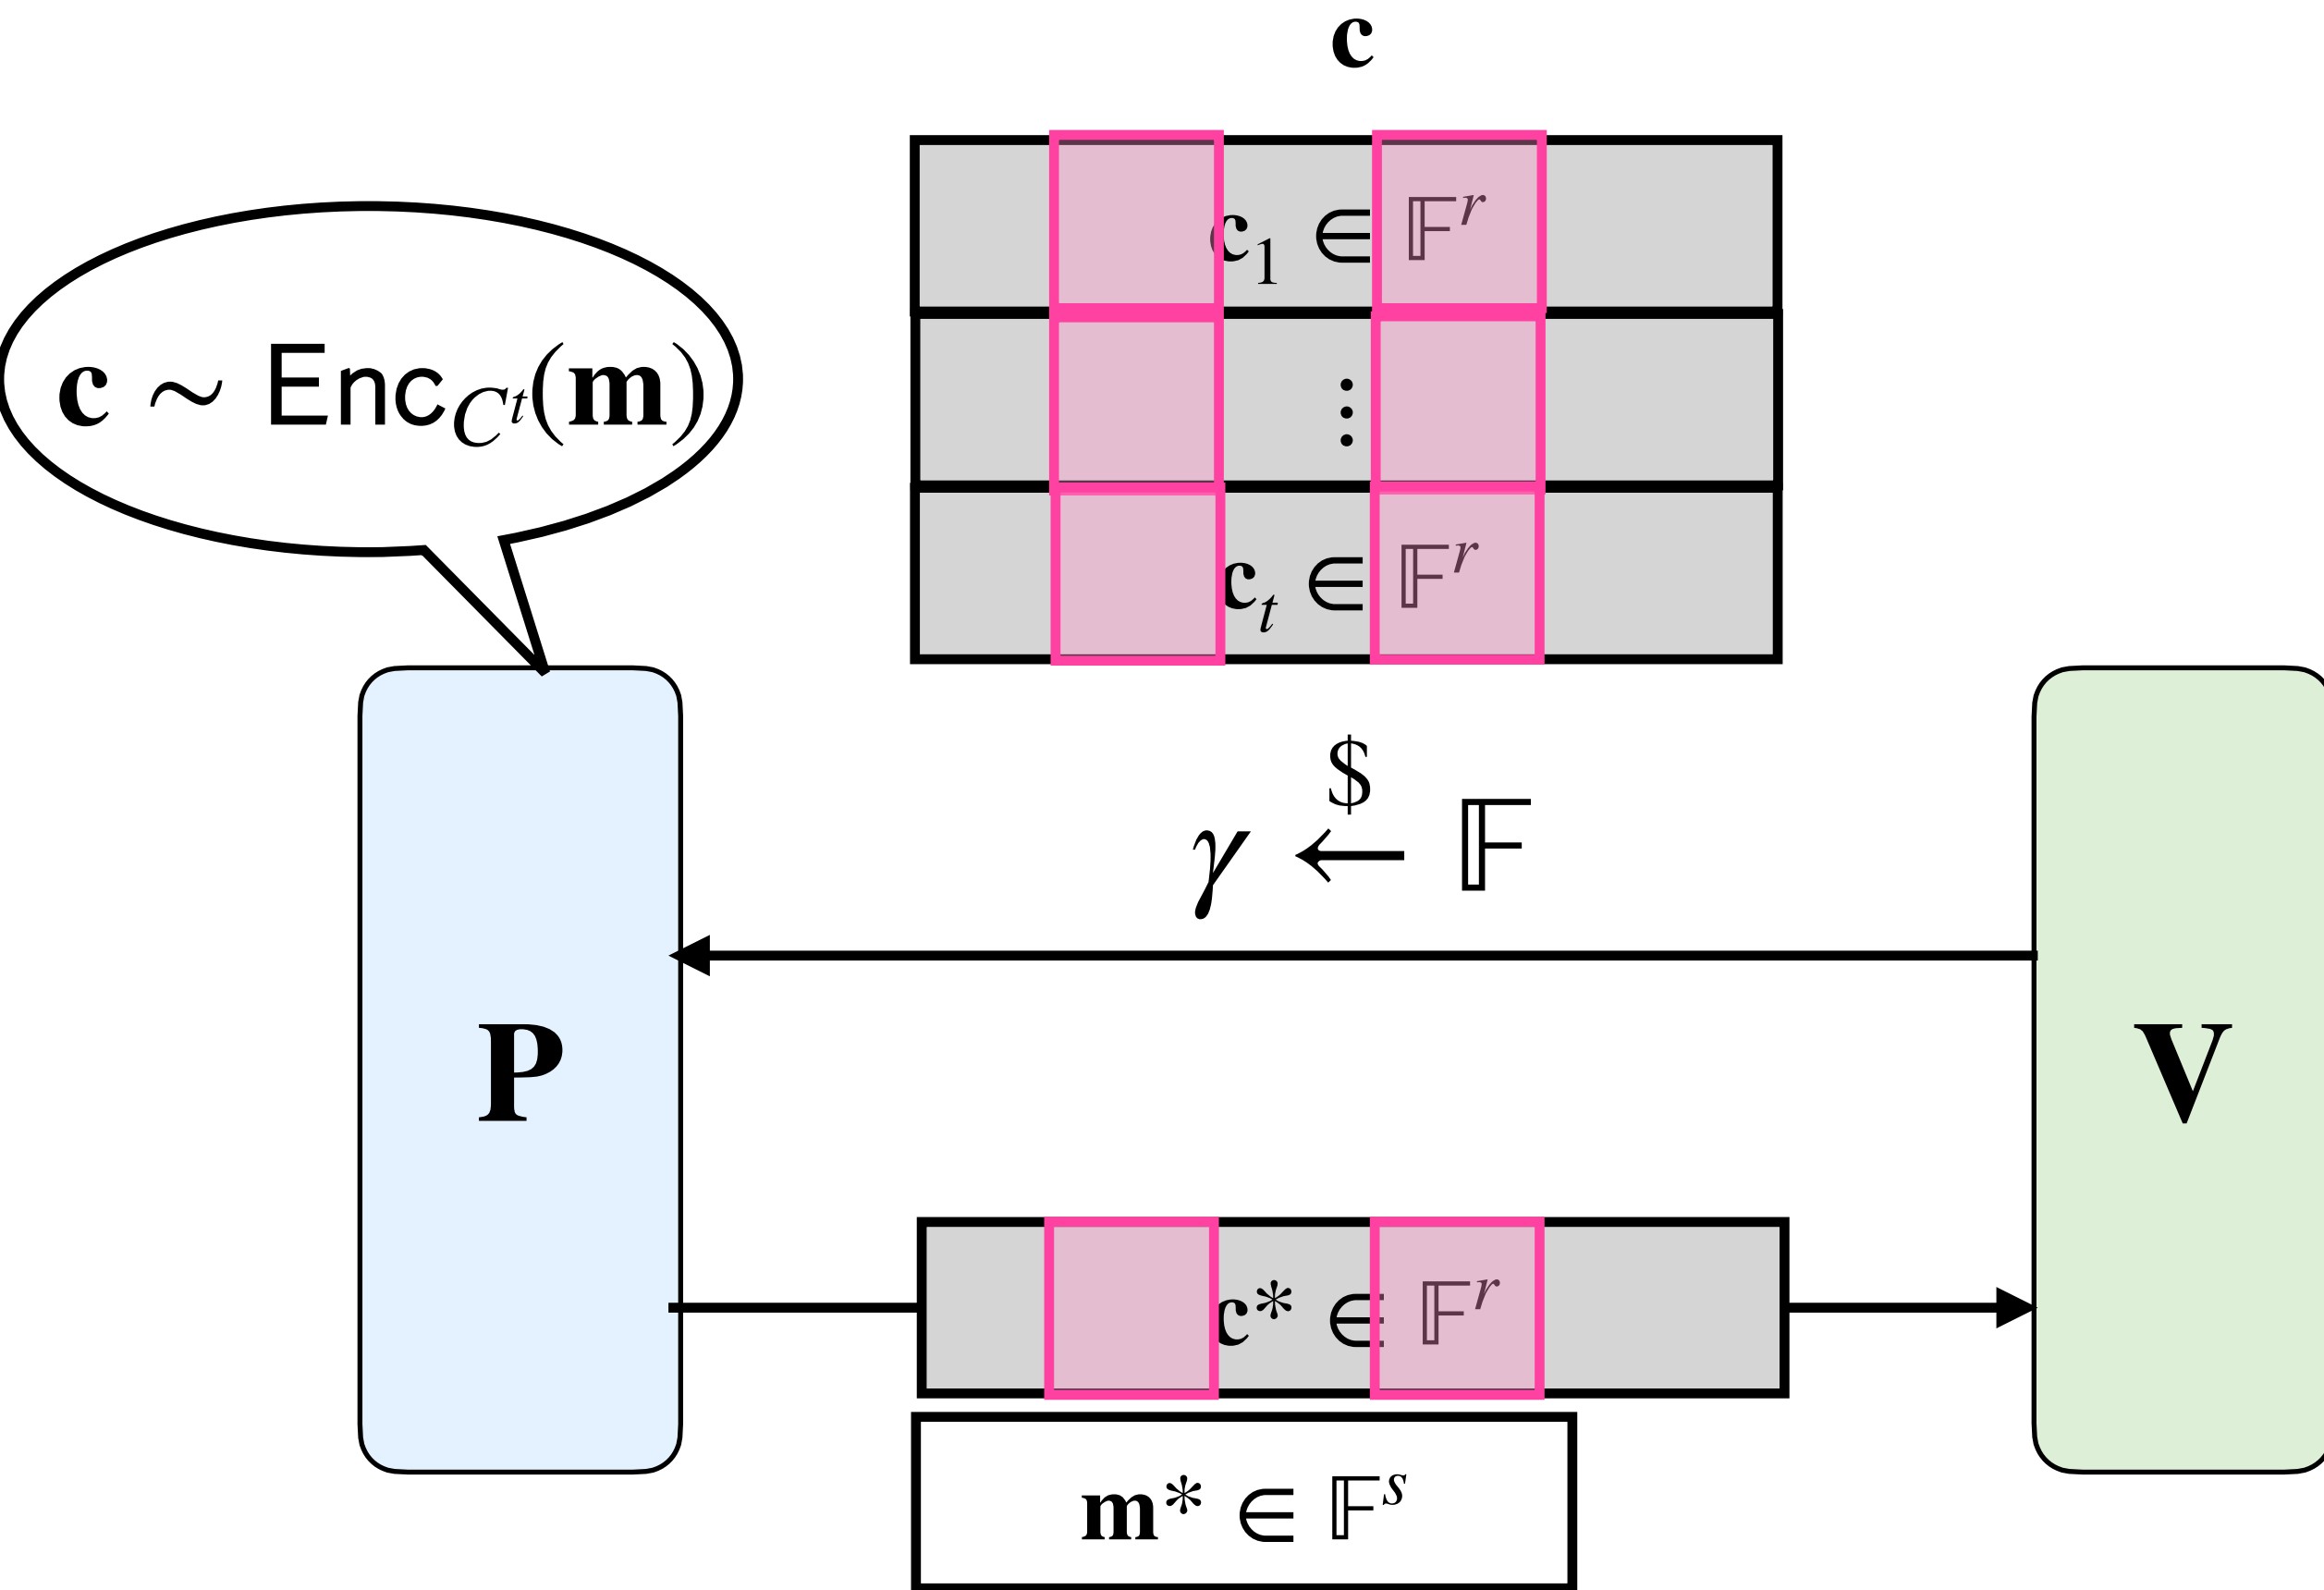
Encoding time $T_{C^t} = t \cdot T_C$

Distance $\delta_{C^t} = \delta_C/t$

IOPP for Interleaved Codes [AHIV17]




IOPP for Interleaved Codes [AHIV17]



1. V sends $\gamma \leftarrow \mathbb{F}$
2. P sends \mathbf{c}^* that is supposed to be $\sum_i \gamma^i \mathbf{c}_i$ and \mathbf{m}^* such that $\mathbf{c}^* = \text{Enc}_C(\mathbf{m}^*)$
3. V checks $\mathbf{c}^* = \sum_i \gamma^i \mathbf{c}_i$ via $\nu = O(\lambda)$ **spot checks:**
 1. V samples ν indices j_1, \dots, j_ν .
 2. For each index j :
 1. V queries $\mathbf{c}_1[j], \dots, \mathbf{c}_t[j]$ and $\mathbf{c}^*[j]$
 2. V checks that $\mathbf{c}^*[j] = \sum_i \gamma^i \cdot \mathbf{c}_i[j]$
4. V checks that \mathbf{c}^* is close to a codeword in C by checking that $\mathbf{c}^* = \text{Enc}_C(\mathbf{m}^*)$

IOPP for Interleaved Codes [AHIV17]



Completeness:

\mathbf{P} is honest, so $\mathbf{c}^* = \sum_i \gamma^i \mathbf{c}_i$, so spot checks pass

Honest \mathbf{P} can set $\mathbf{m}^* = \sum_i \gamma^i \mathbf{m}_i$

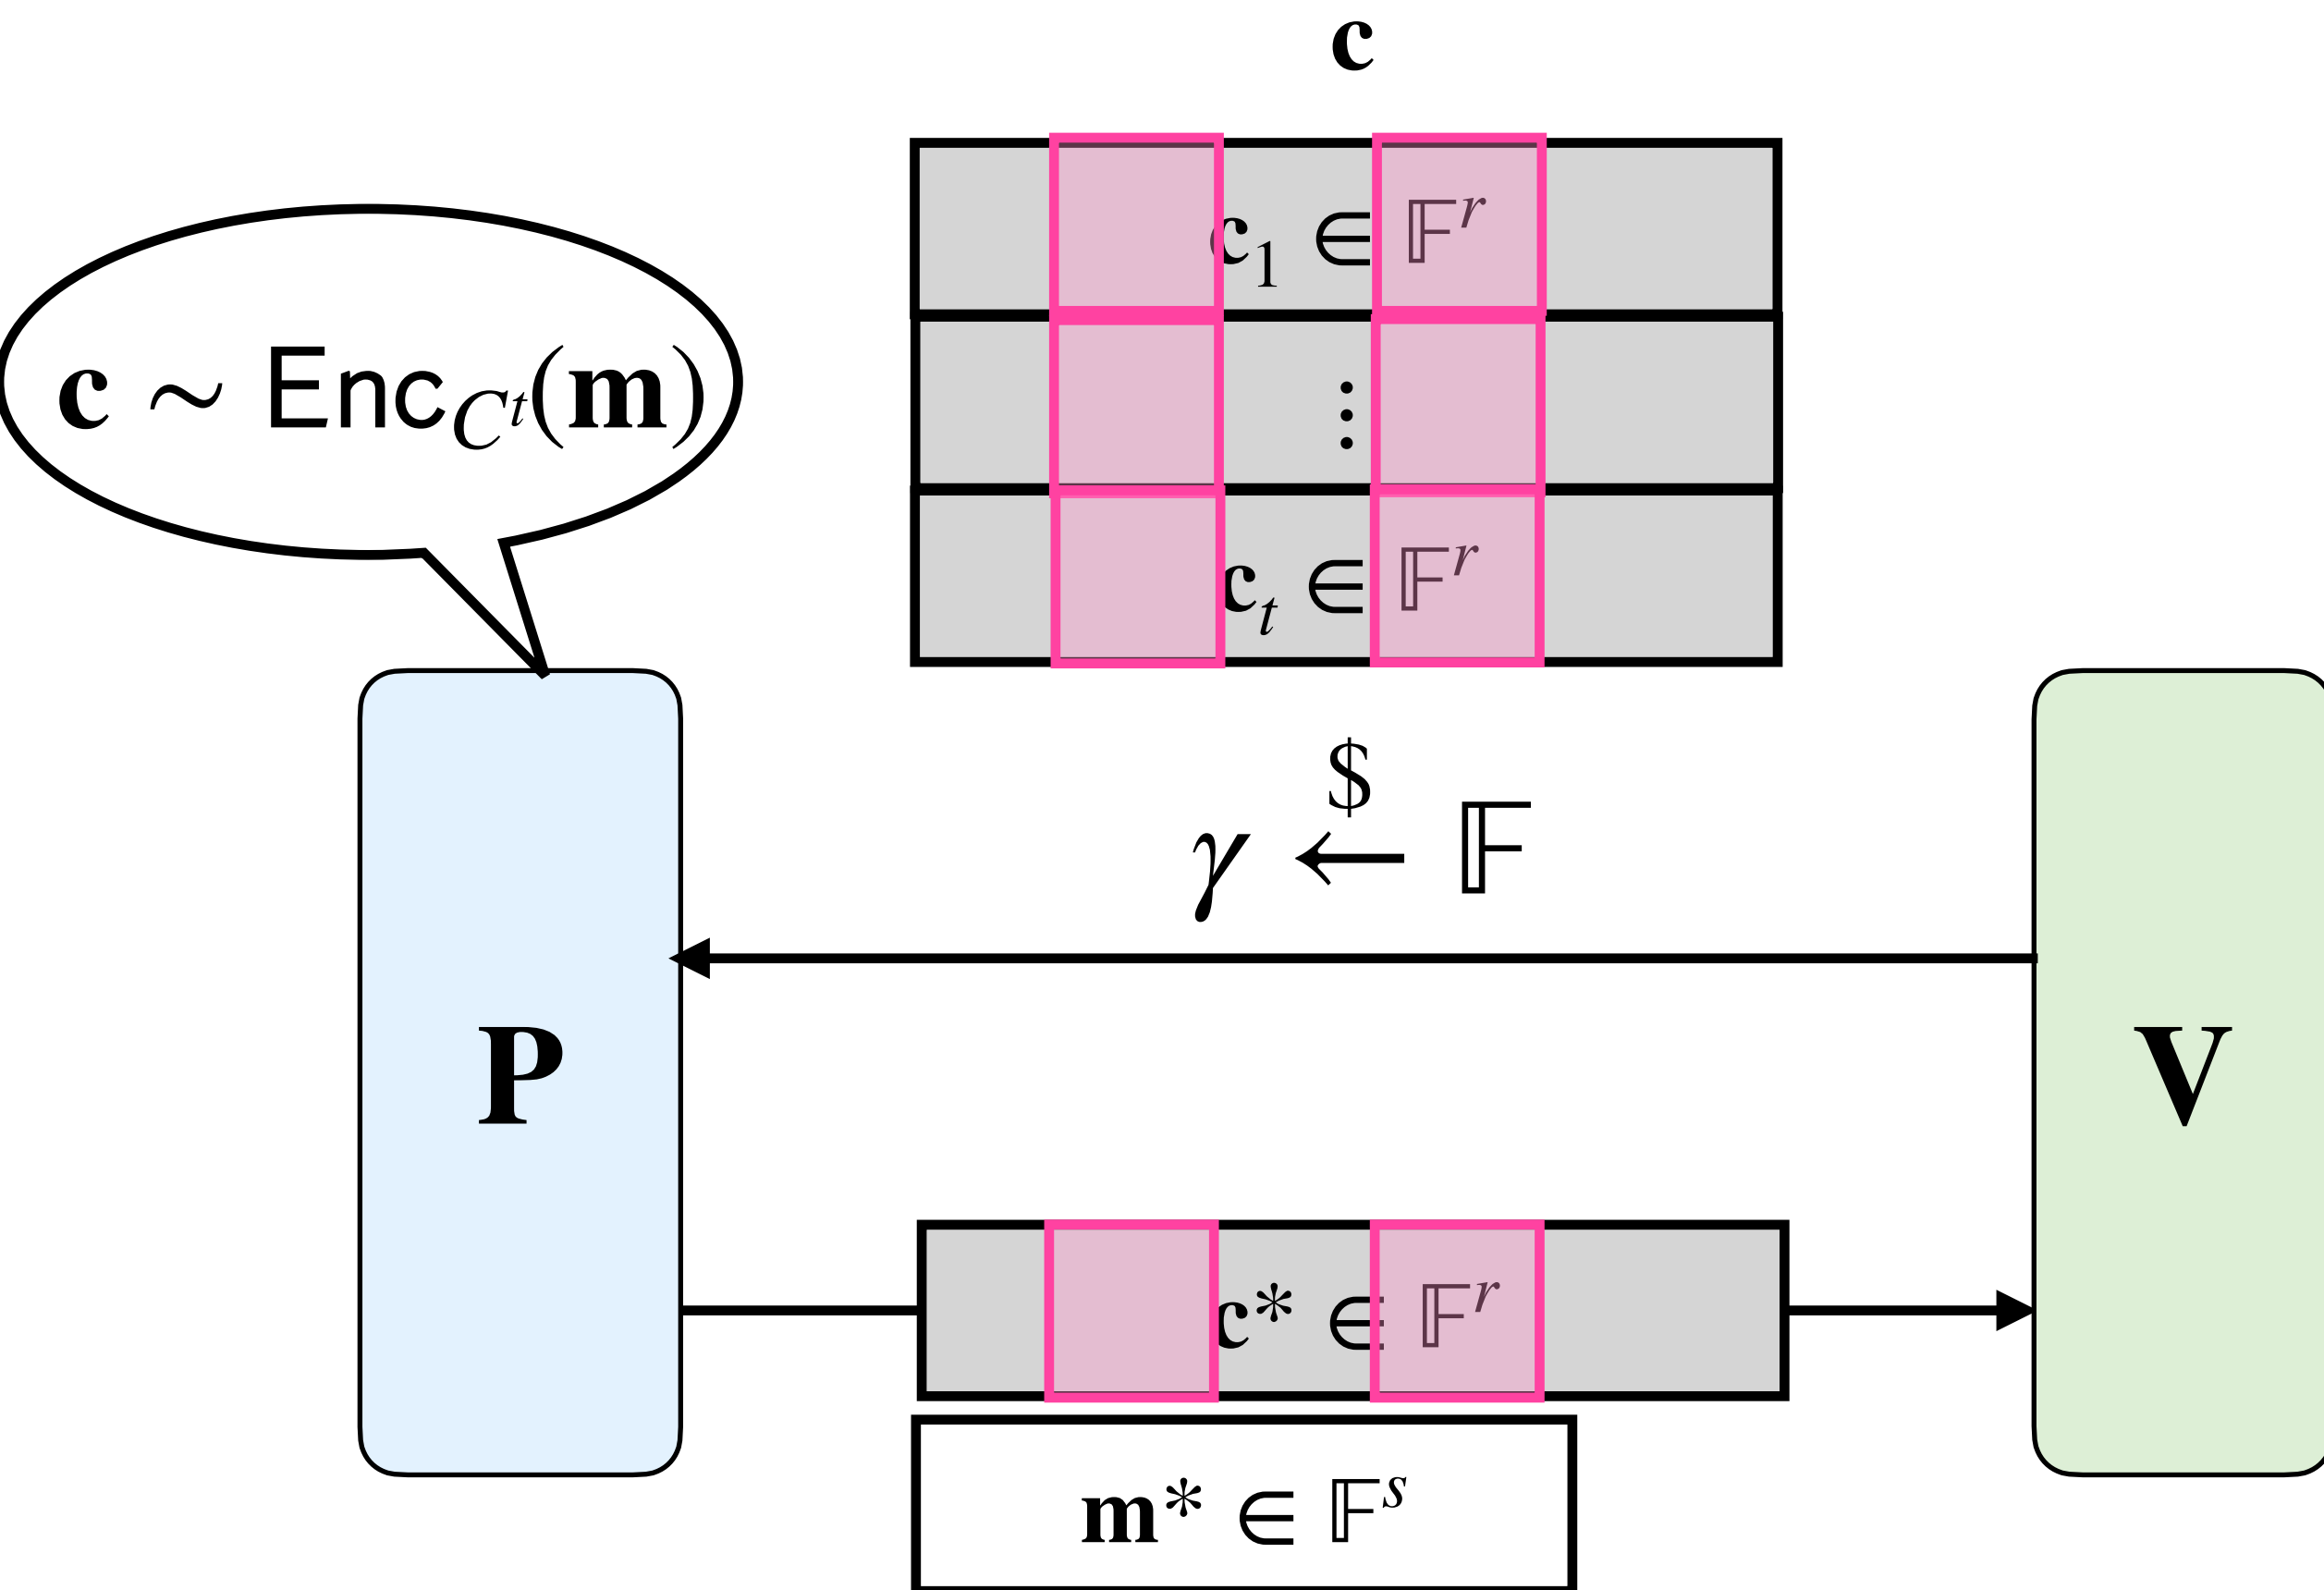
Since C is linear, $\text{Enc}_C(\mathbf{m}^*) = \text{Enc}_C\left(\sum_i \gamma^i \mathbf{m}_i\right)$

$$= \sum_i \gamma^i \text{Enc}_C(\mathbf{m}_i)$$

$$= \sum_i \gamma^i \mathbf{c}_i$$

$$= \mathbf{c}^*$$

IOPP for Interleaved Codes [AHIV17]



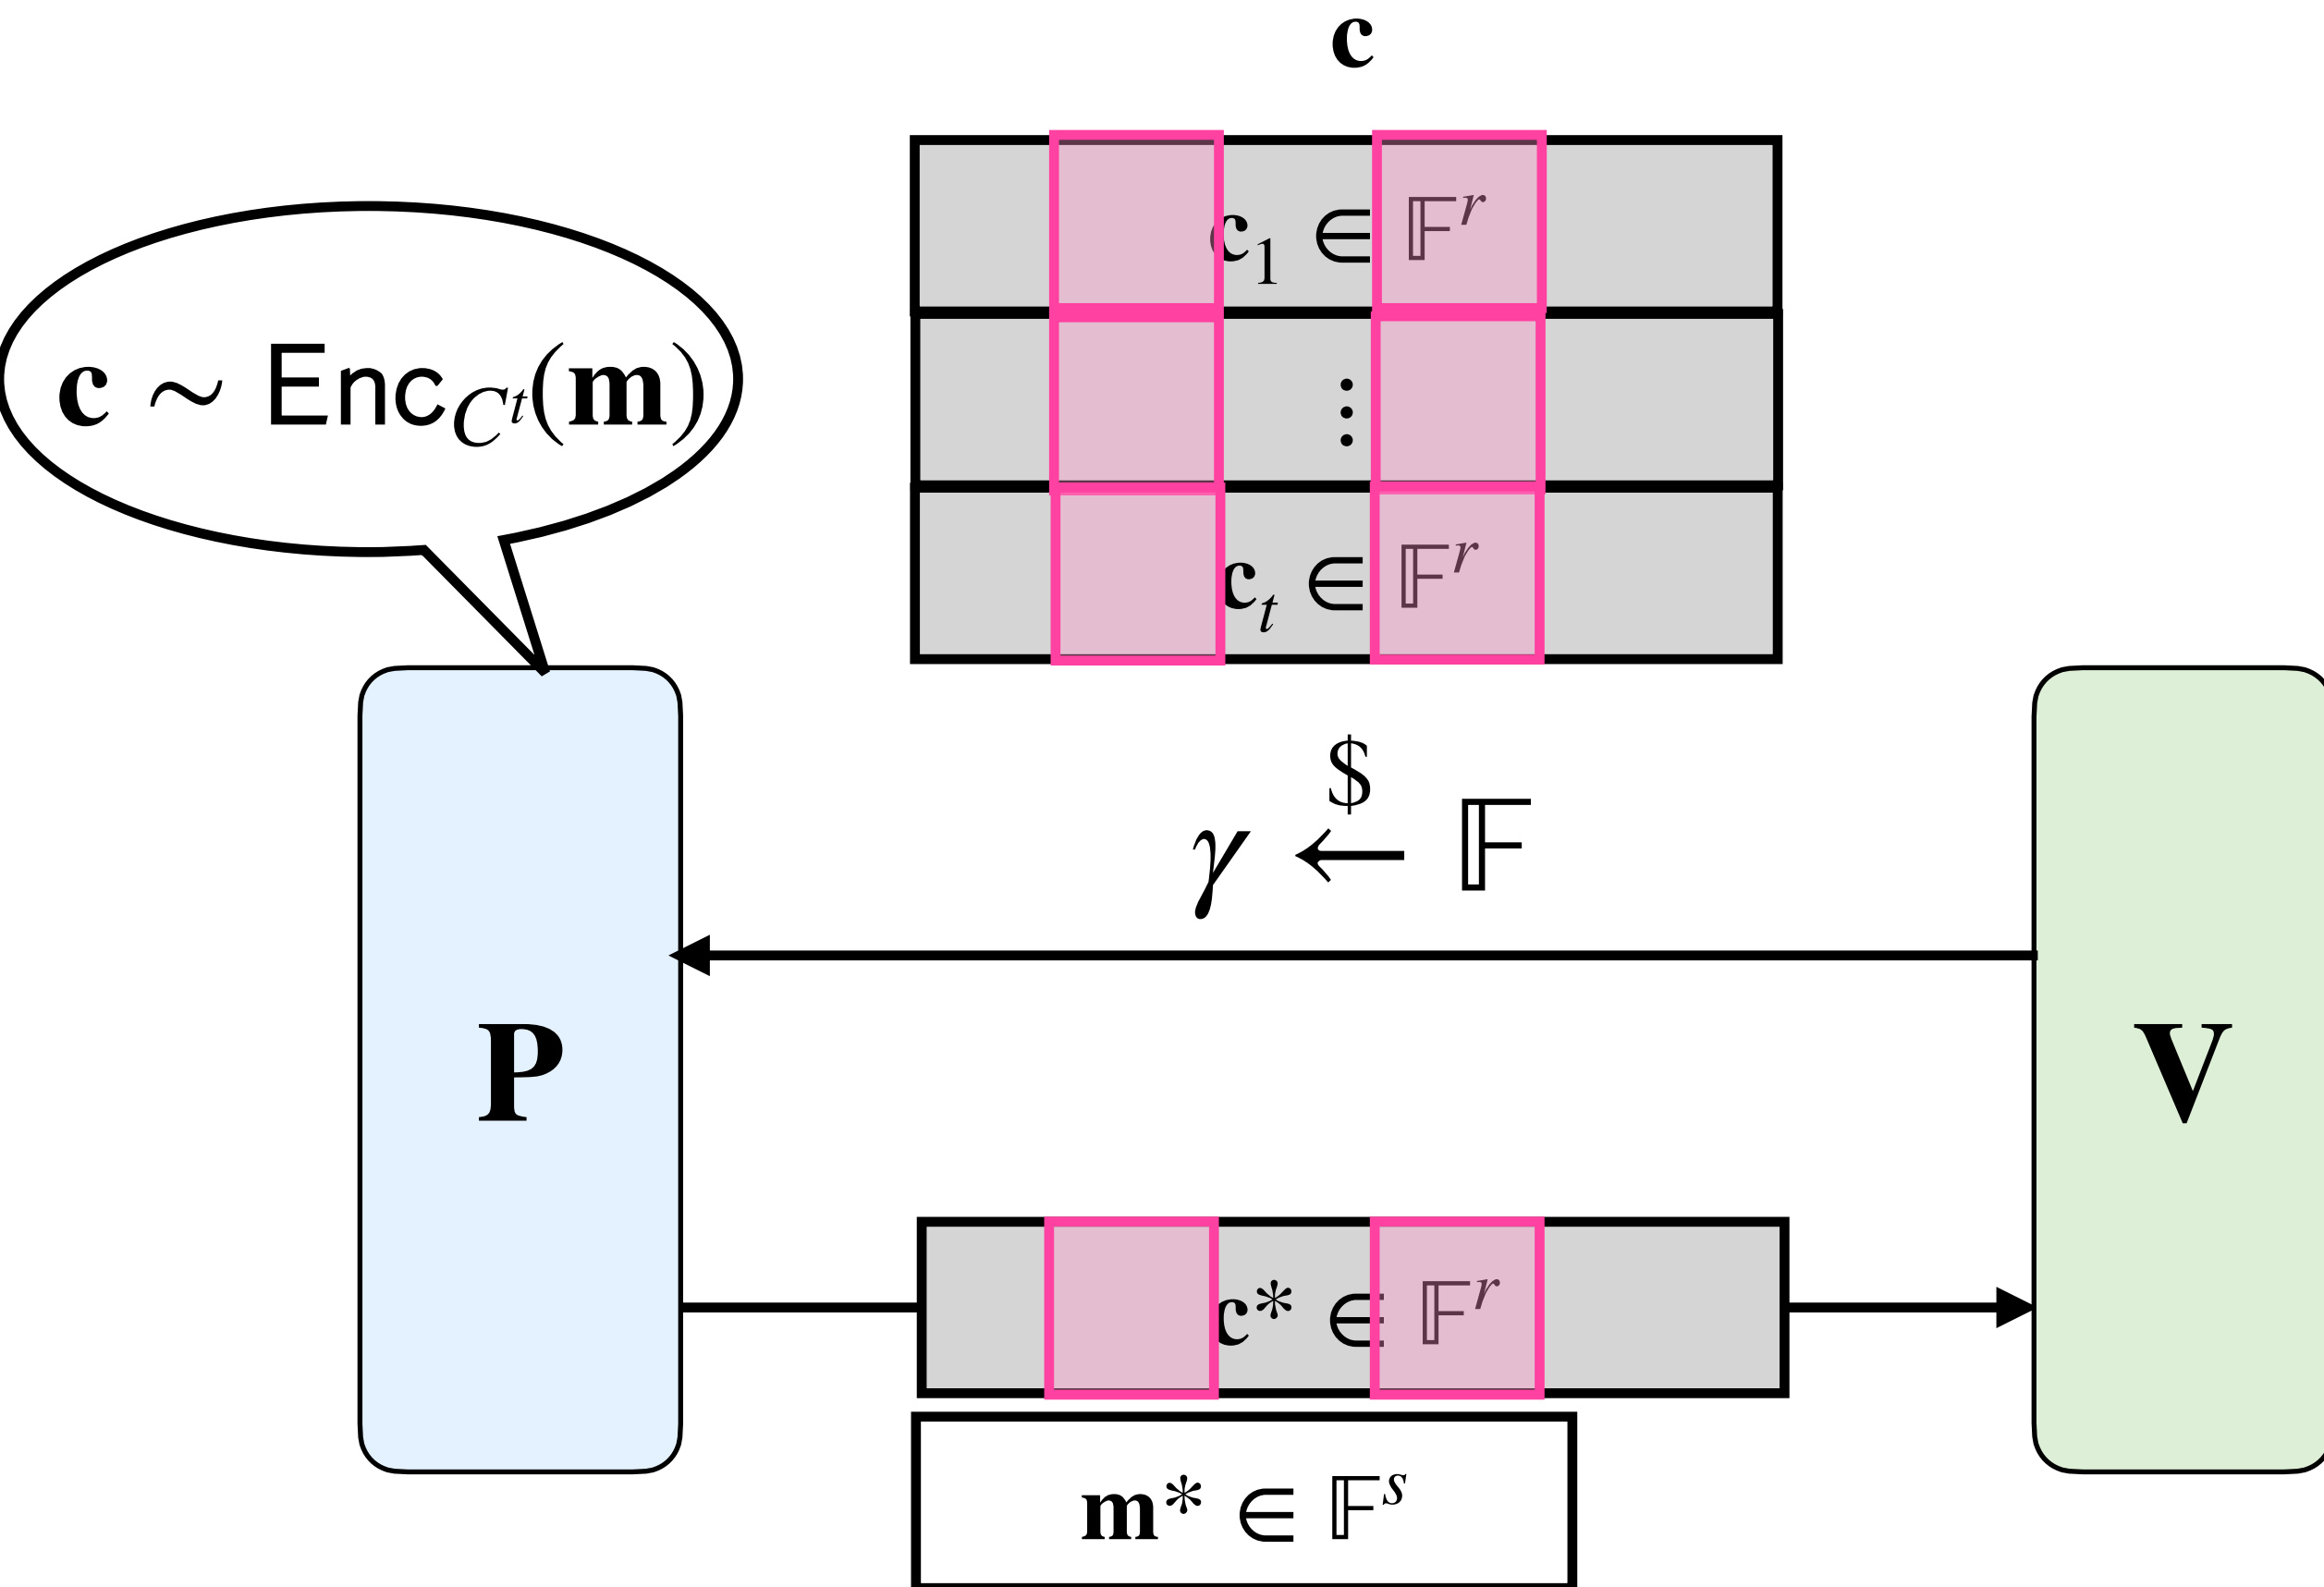
Soundness: \mathbf{V} accepts but at least one \mathbf{c}_i is far from \mathcal{C}

Case 0: \mathbf{c}^* is far from \mathcal{C}
 \mathbf{V} rejects in re-encoding check

Case 1.1: \mathbf{c}^* is close to \mathcal{C} and $\mathbf{c}^* = \sum_i \gamma^i \mathbf{c}_i$
 Happens with negligible probability due to proximity gaps [Dan's talk!].

Case 1.2: \mathbf{c}^* is close to \mathcal{C} and far from $\sum_i \gamma^i \mathbf{c}_i$
 Spot checks all pass with probability $(1 - \delta_C)^\nu$;
 set $\nu = -\lambda / \log_2(1 - \delta_C)$ to make this negligible.

IOPP for Interleaved Codes [AHIV17]



Efficiency:

prover time: $T_C + O(k)$

queries: $\nu \cdot t + s \mathbb{F}$.

when encoding k -length messages, this is minimized when $t = \sqrt{k/\nu}$ and $s = \sqrt{k\nu}$, resulting in query complexity of

$$O(\sqrt{k\nu})$$


Chapter 2

Query Reduction by Interleaving Less

Goal:

IOPP with $O(\lambda \log k)$ queries

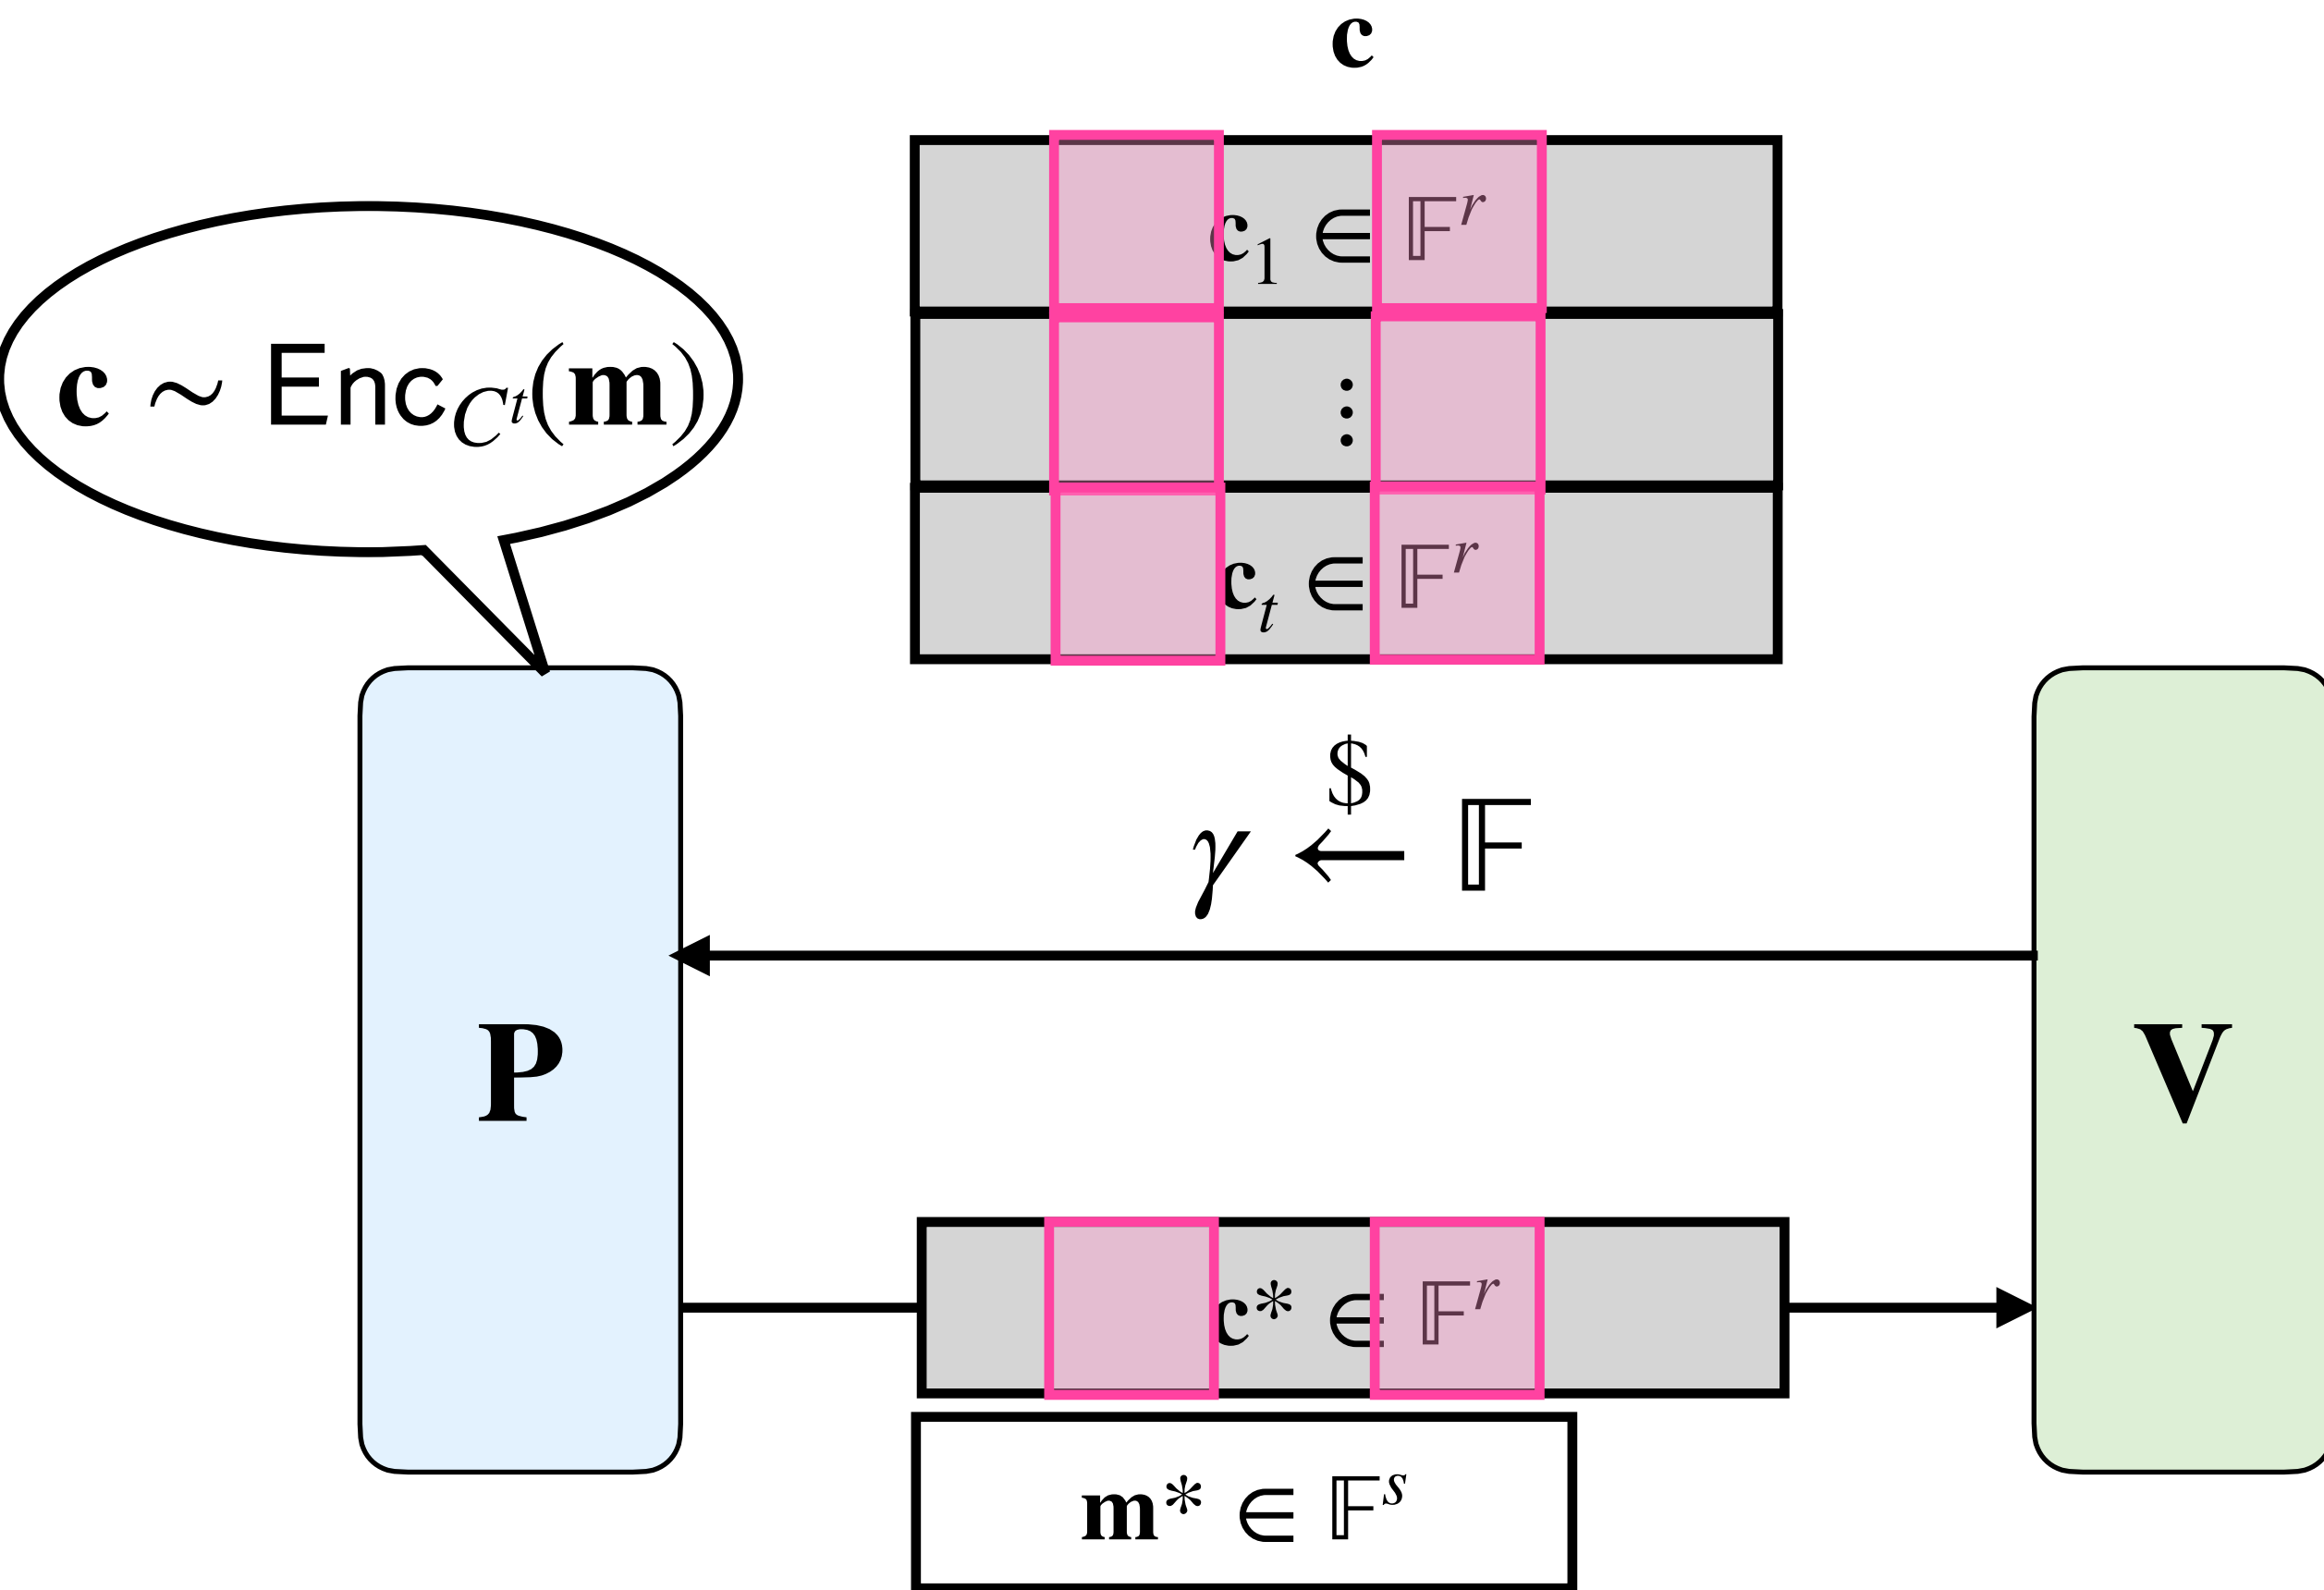
Why is the IOPP Expensive?



1. V sends $\gamma \leftarrow \mathbb{F}$
2. P sends \mathbf{c}^* that is supposed to be $\sum_i \gamma^i \mathbf{c}_i$ and \mathbf{m}^* such that $\mathbf{c}^* = \text{Enc}_C(\mathbf{m}^*)$
3. V checks $\mathbf{c}^* = \sum_i \gamma^i \mathbf{c}_i$ via $\nu = O(\lambda)$ **spot checks:**
 1. V samples ν indices j_1, \dots, j_ν .
 2. For each index j :
 1. V queries $\mathbf{c}_1[j], \dots, \mathbf{c}_t[j]$ and $\mathbf{c}^*[j]$
 2. V checks that $\mathbf{c}^*[j] = \sum_i \gamma^i \cdot \mathbf{c}_i[j]$

4. V checks that \mathbf{c}^* is close to a codeword in C by checking that $\mathbf{c}^* = \text{Enc}_C(\mathbf{m}^*)$

Why is the IOPP Expensive?



1. V sends $\gamma \xleftarrow{\$} \mathbb{F}$
2. P sends \mathbf{c}^* that is supposed to be $\sum_i \gamma^i \mathbf{c}_i$ and \mathbf{m}^* such that $\mathbf{c}^* = \text{Enc}_C(\mathbf{m}^*)$
3. V checks $\mathbf{c}^* = \sum_i \gamma^i \mathbf{c}_i$ via $\nu = O(\lambda)$ **spot checks:**
 1. V samples ν indices j_1, \dots, j_ν .
 2. For each index j :
 1. V queries $\mathbf{c}_1[j], \dots, \mathbf{c}_t[j]$ and $\mathbf{c}^*[j]$
 2. V checks that $\mathbf{c}^*[j] = \sum_i \gamma^i \cdot \mathbf{c}_i[j]$

4. V checks that \mathbf{c}^* is close to a codeword in C by checking that $\mathbf{c}^* = \text{Enc}_C(\mathbf{m}^*)$

Why is the IOPP Expensive?

1. \mathbf{V} sends $\gamma \xrightarrow{\$} \mathbb{F}$
2. \mathbf{P} sends \mathbf{c}^* that is supposed to be $\sum_i \gamma^i \mathbf{c}_i$ and \mathbf{m}^* such that $\mathbf{c}^* = \text{Enc}_C(\mathbf{m}^*)$
3. \mathbf{V} checks $\mathbf{c}^* = \sum_i \gamma^i \mathbf{c}_i$ via $\nu = O(\lambda)$
spot checks:
 1. \mathbf{V} samples ν indices j_1, \dots, j_ν .
 2. For each index j :
 1. \mathbf{V} queries $\mathbf{c}_1[j], \dots, \mathbf{c}_t[j]$ and $\mathbf{c}^*[j]$
 2. \mathbf{V} checks that $\mathbf{c}^*[j] = \sum_i \gamma^i \cdot \mathbf{c}_i[j]$

4. \mathbf{V} checks that \mathbf{c}^* is close to a codeword in C by checking that $\mathbf{c}^* = \text{Enc}_C(\mathbf{m}^*)$

Interactive "Reduction":

\mathbf{c} is close to C^t

↓

\mathbf{c}^* is close to C

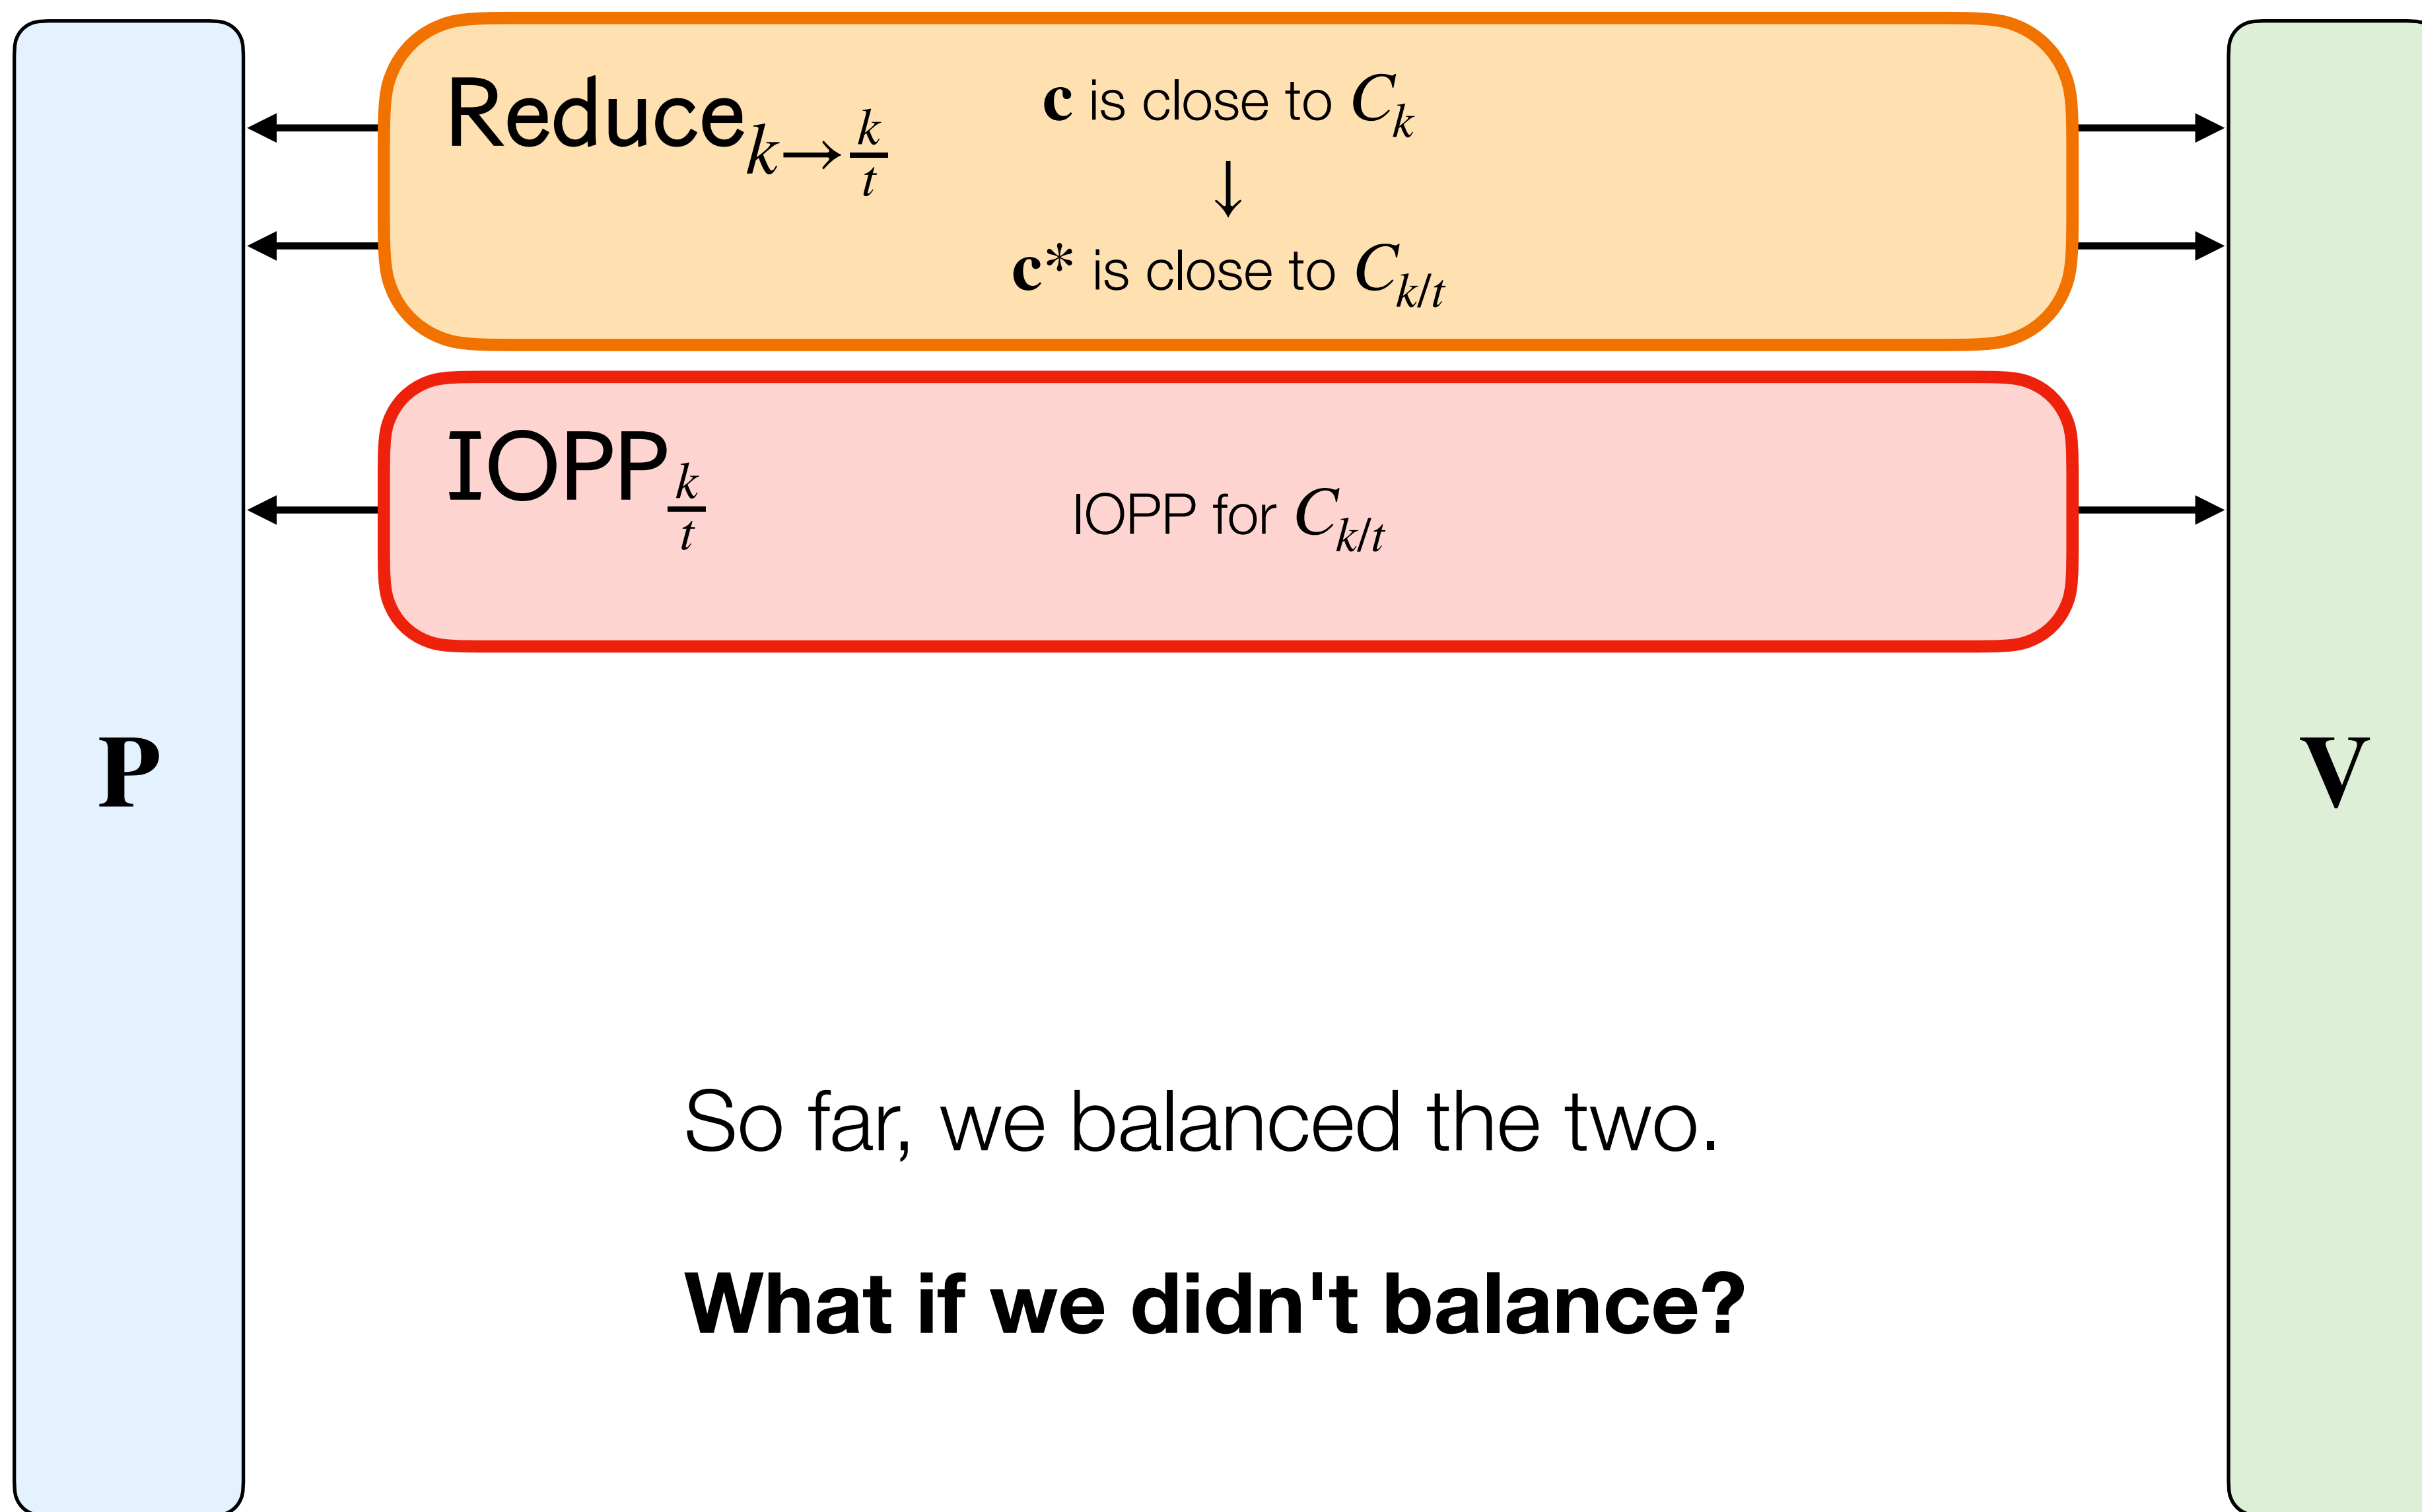
queries: $O(\lambda t)$

Naive IOPP for C

queries: $O(s)$

$$t = \sqrt{k/\lambda}$$

Abstract Protocol

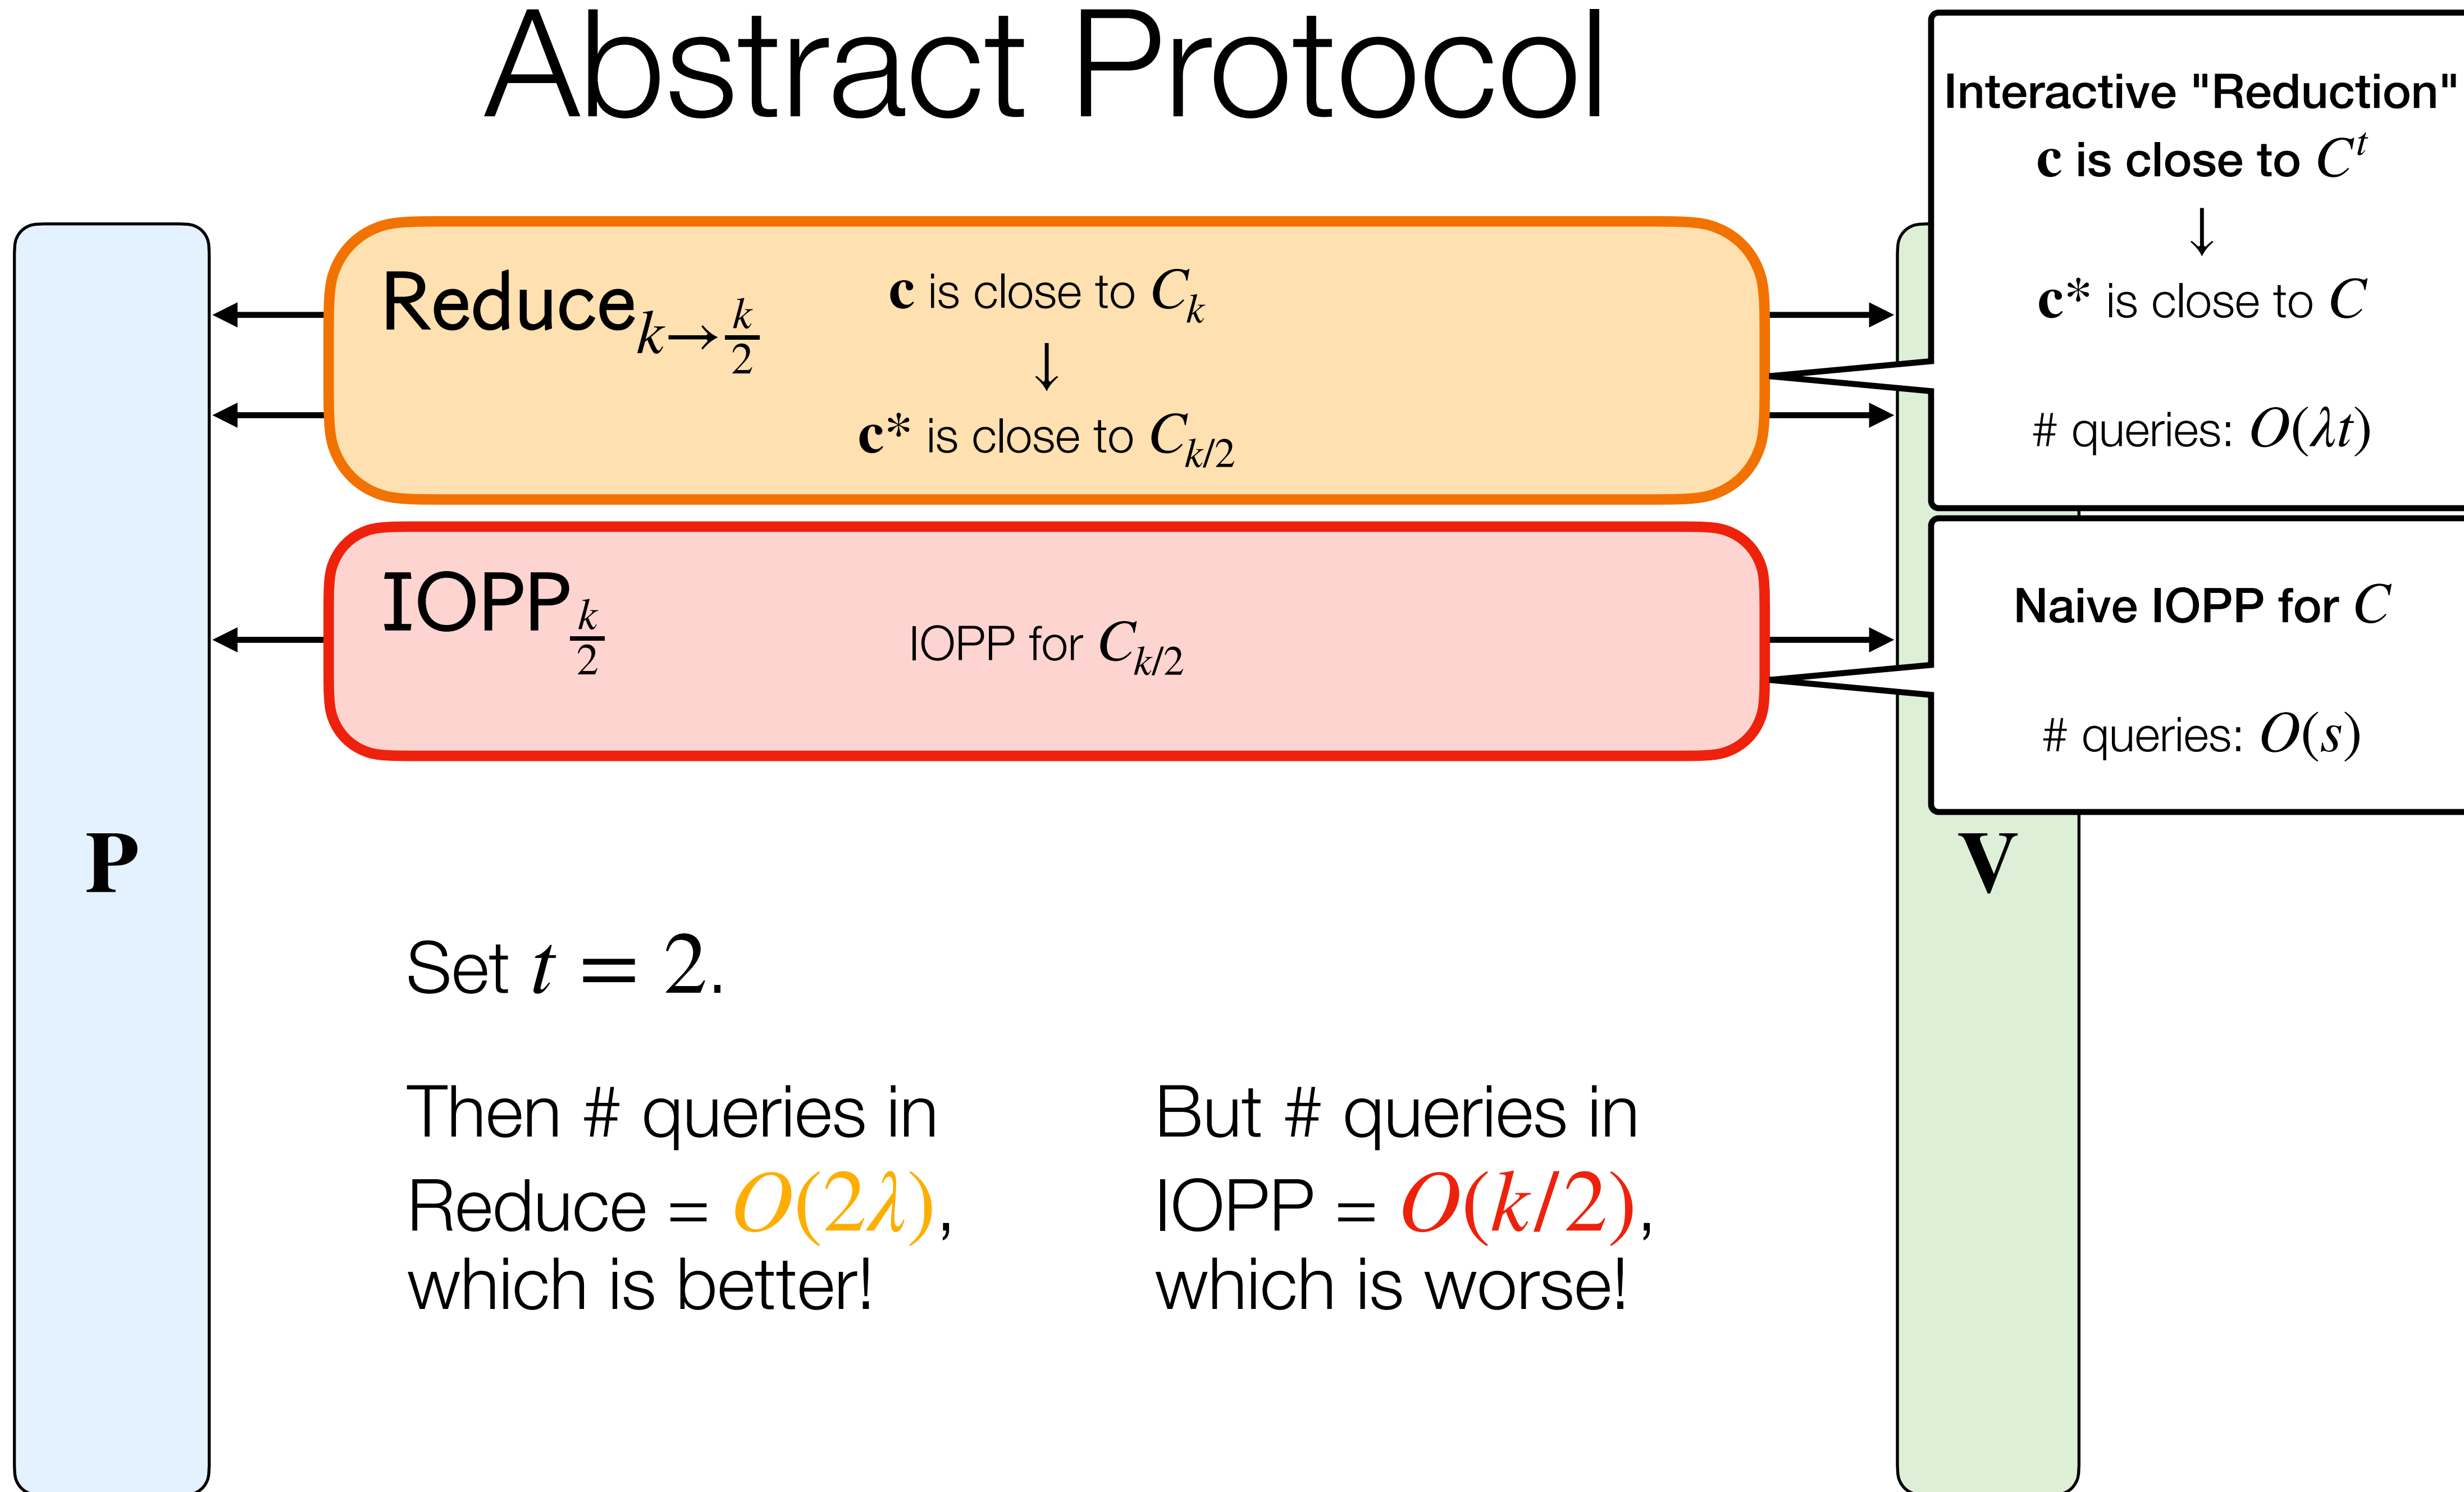


So far, we balanced the two.

What if we didn't balance?

$t = 2$

Abstract Protocol



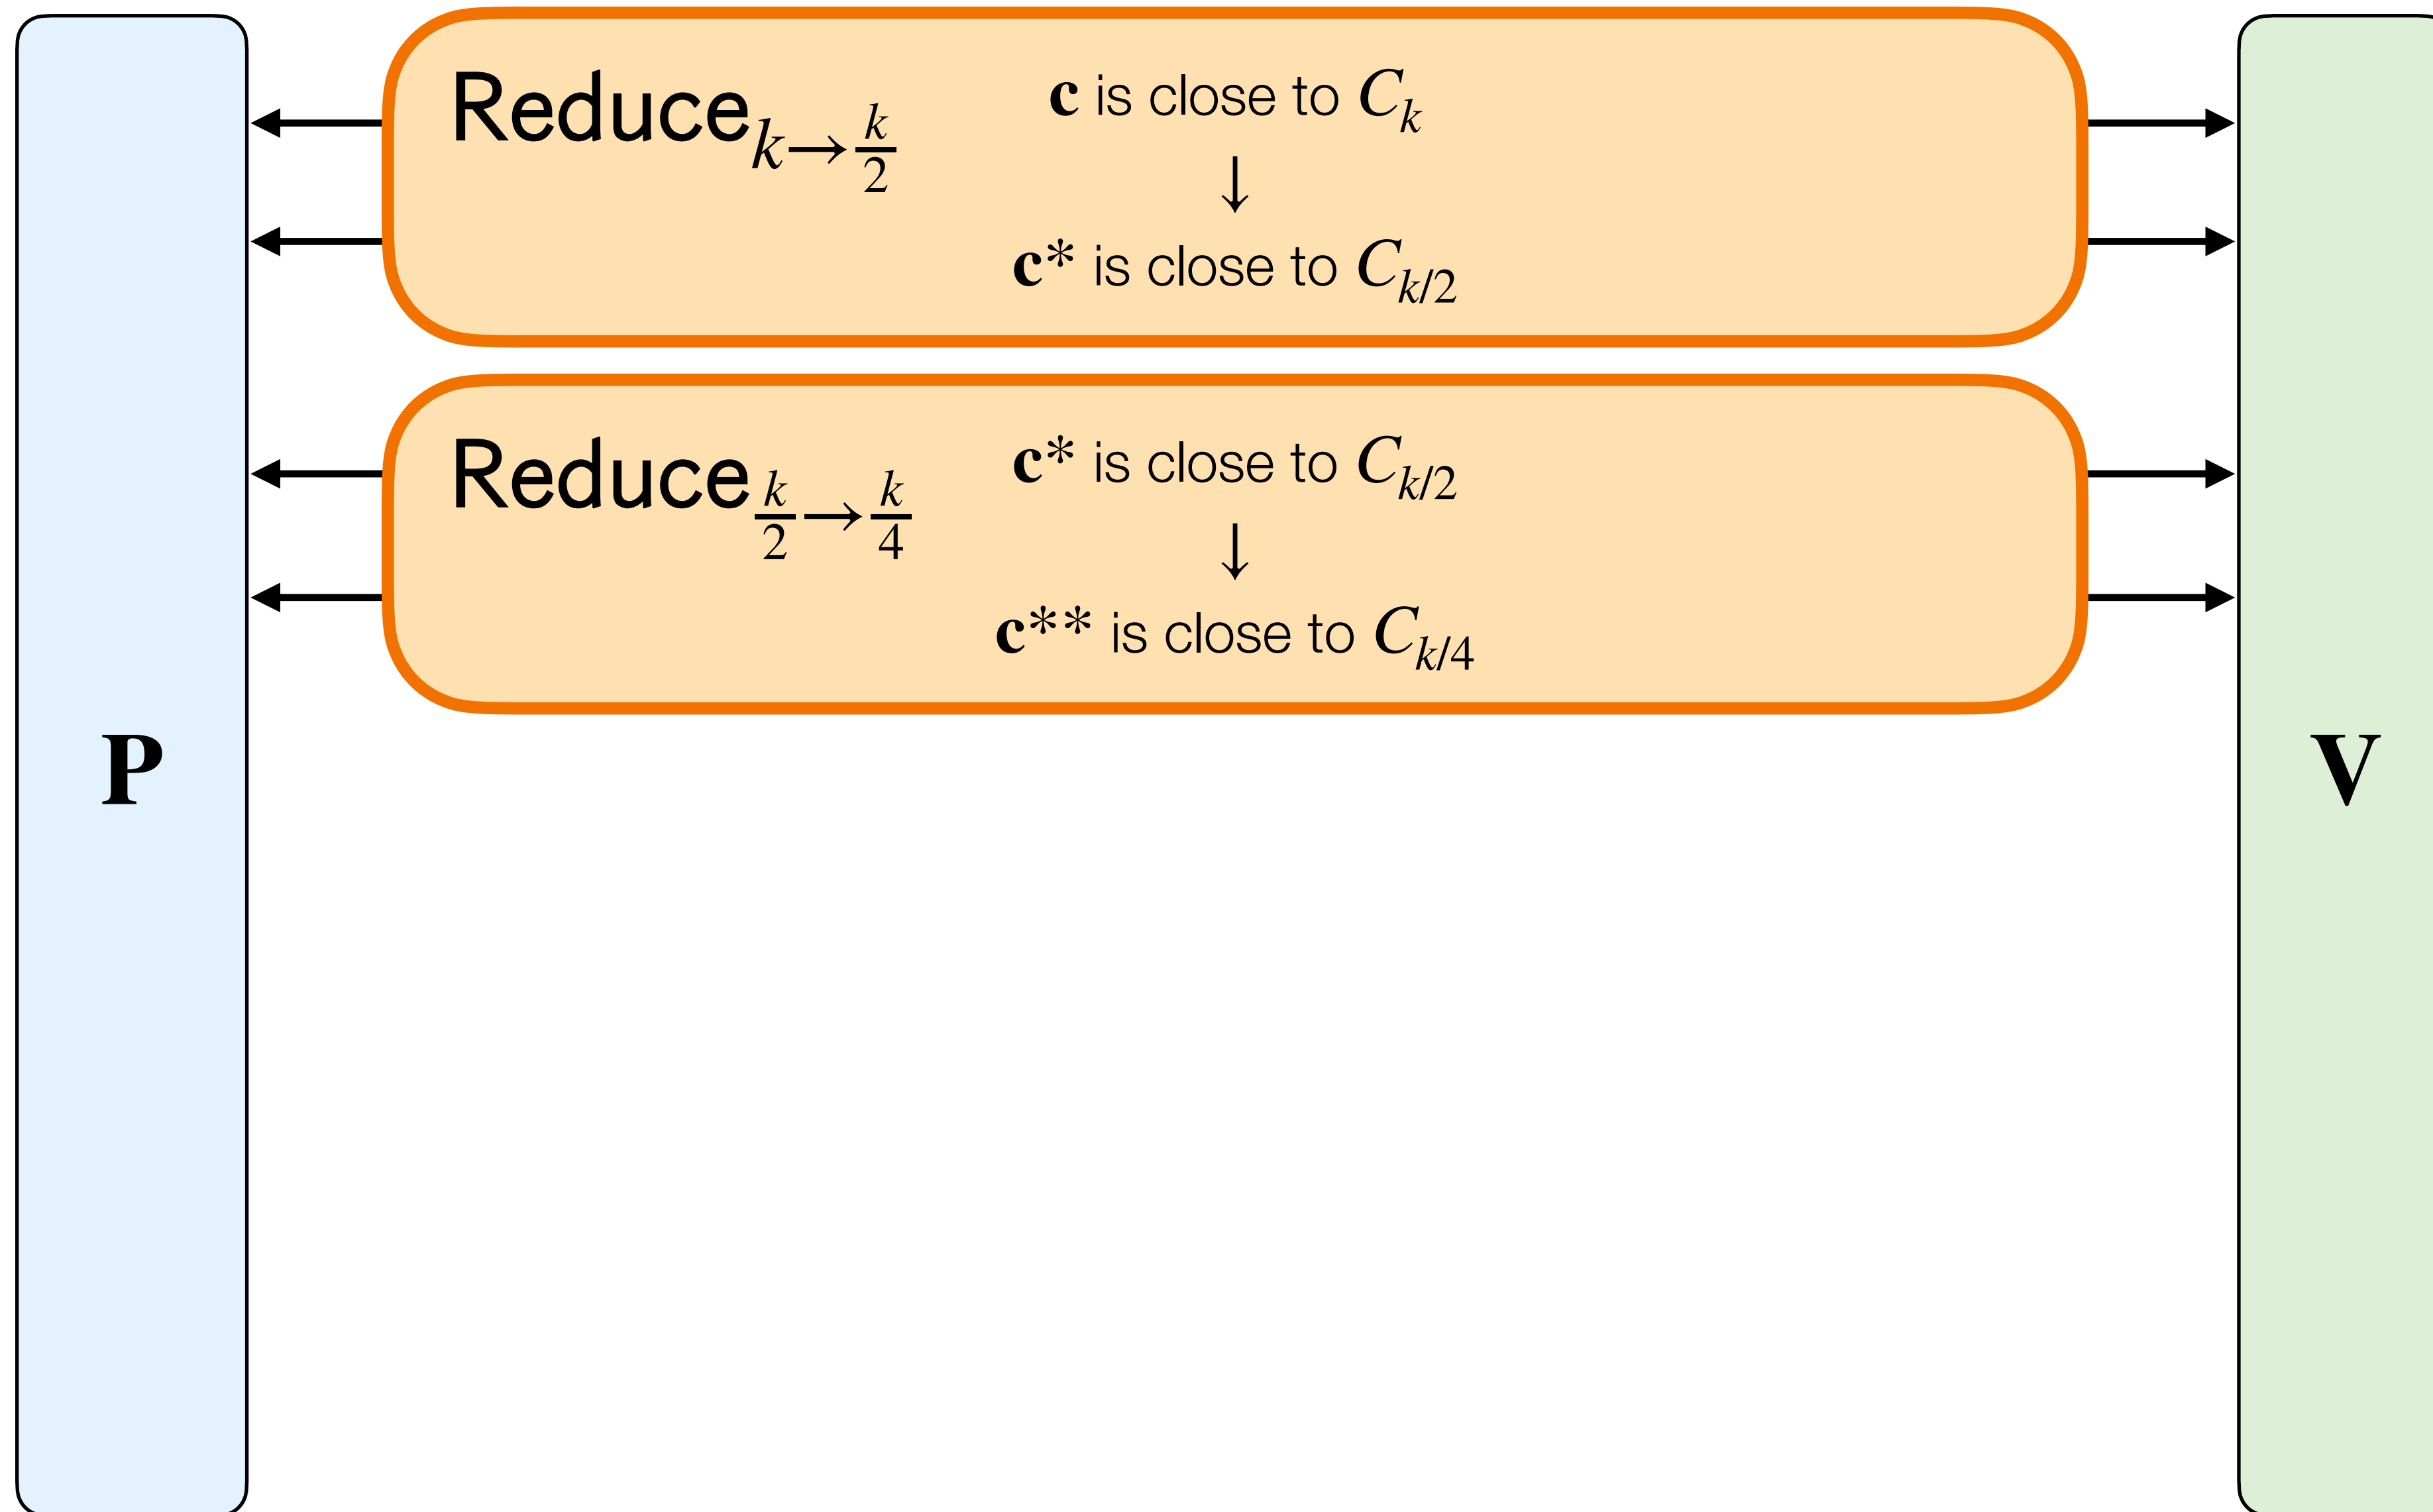
Set $t = 2$.

Then # queries in Reduce = $O(2\lambda)$, which is better!

But # queries in IOPP = $O(k/2)$, which is worse!

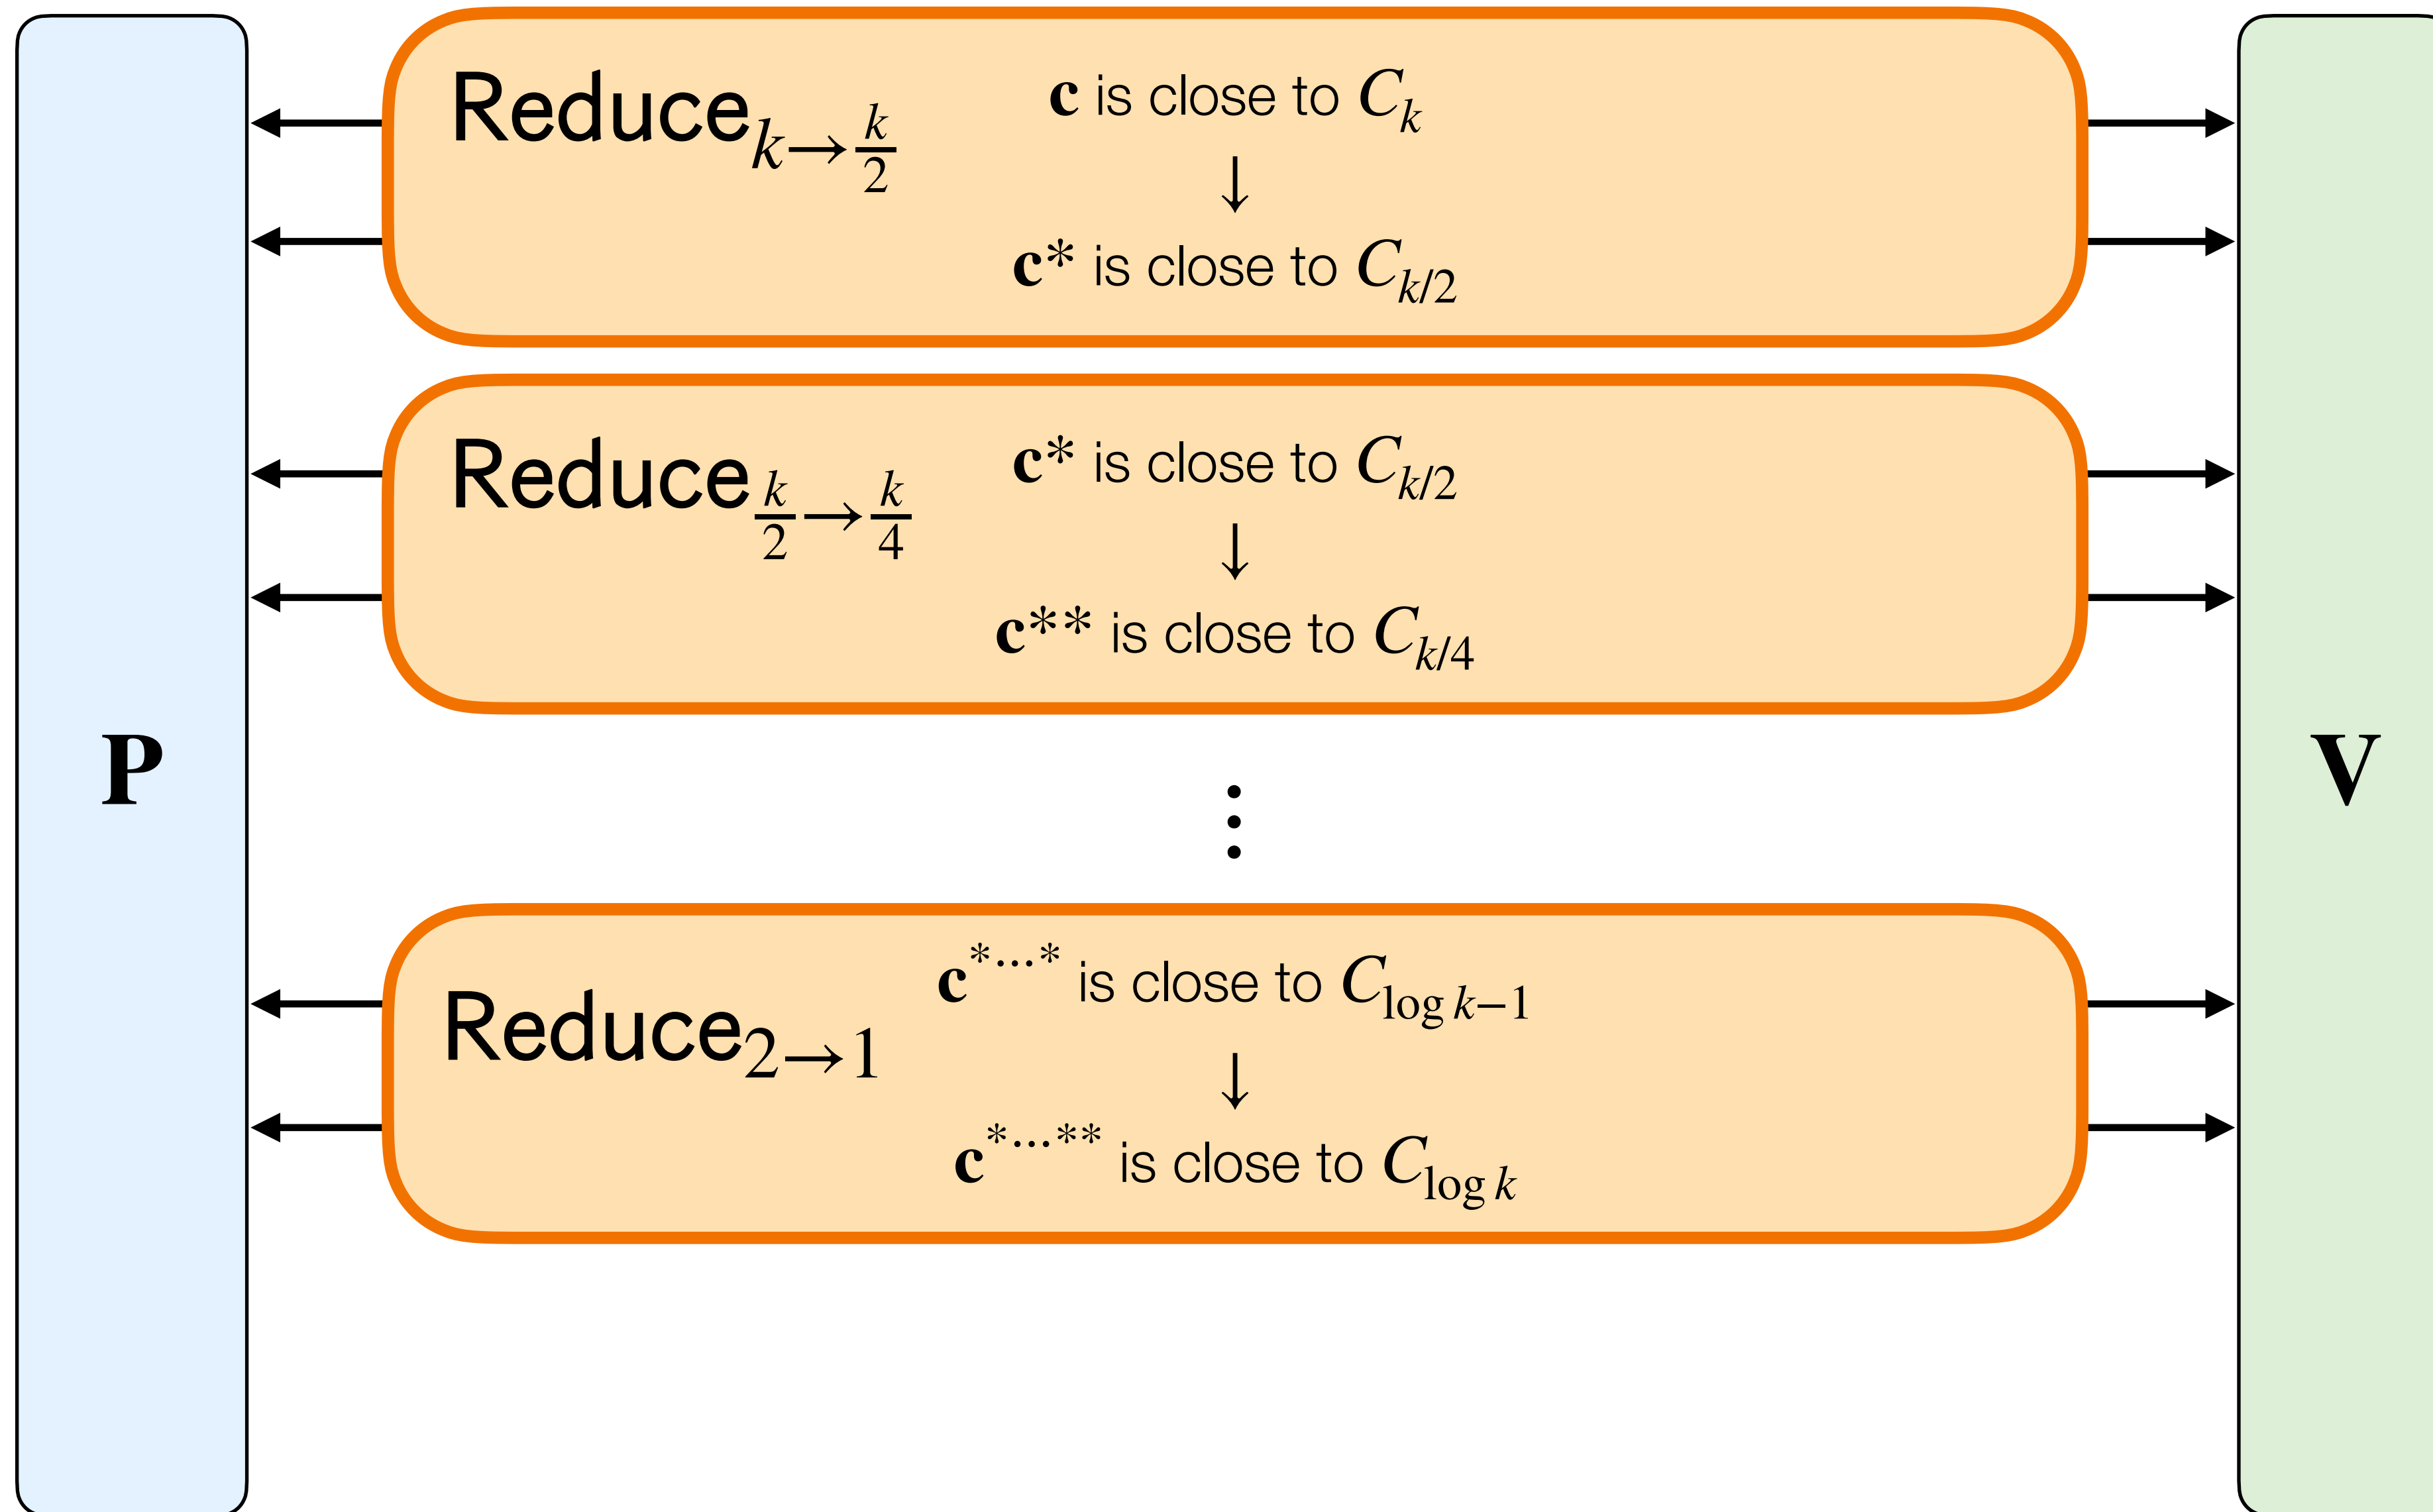
$t = 2$

Recurse!



$t = 2$

Recurse!



Codes and Reductions

Consider the following codes:

$B_{\frac{k}{2}}$ with msg. length $k/2$

$B_{\frac{k}{4}}$ with msg. length $k/4$

⋮

B_1 with msg. length 1

Set $C_i = B_{i/2}^2$

[RR24, RR25, BCFRRZ25, NST25, BMMS26b]

$O(k/2^i)$ prover time and $O(\lambda)$ query IOR for codeswitchable B_i and arbitrary output code.

Reduce $\frac{k}{2^i} \rightarrow \frac{k}{2^{i+1}}$

reduce size

Fold

prox to $C_{\frac{k}{2^i}} \rightarrow$ prox to $B_{\frac{k}{2^{i+1}}}$

+

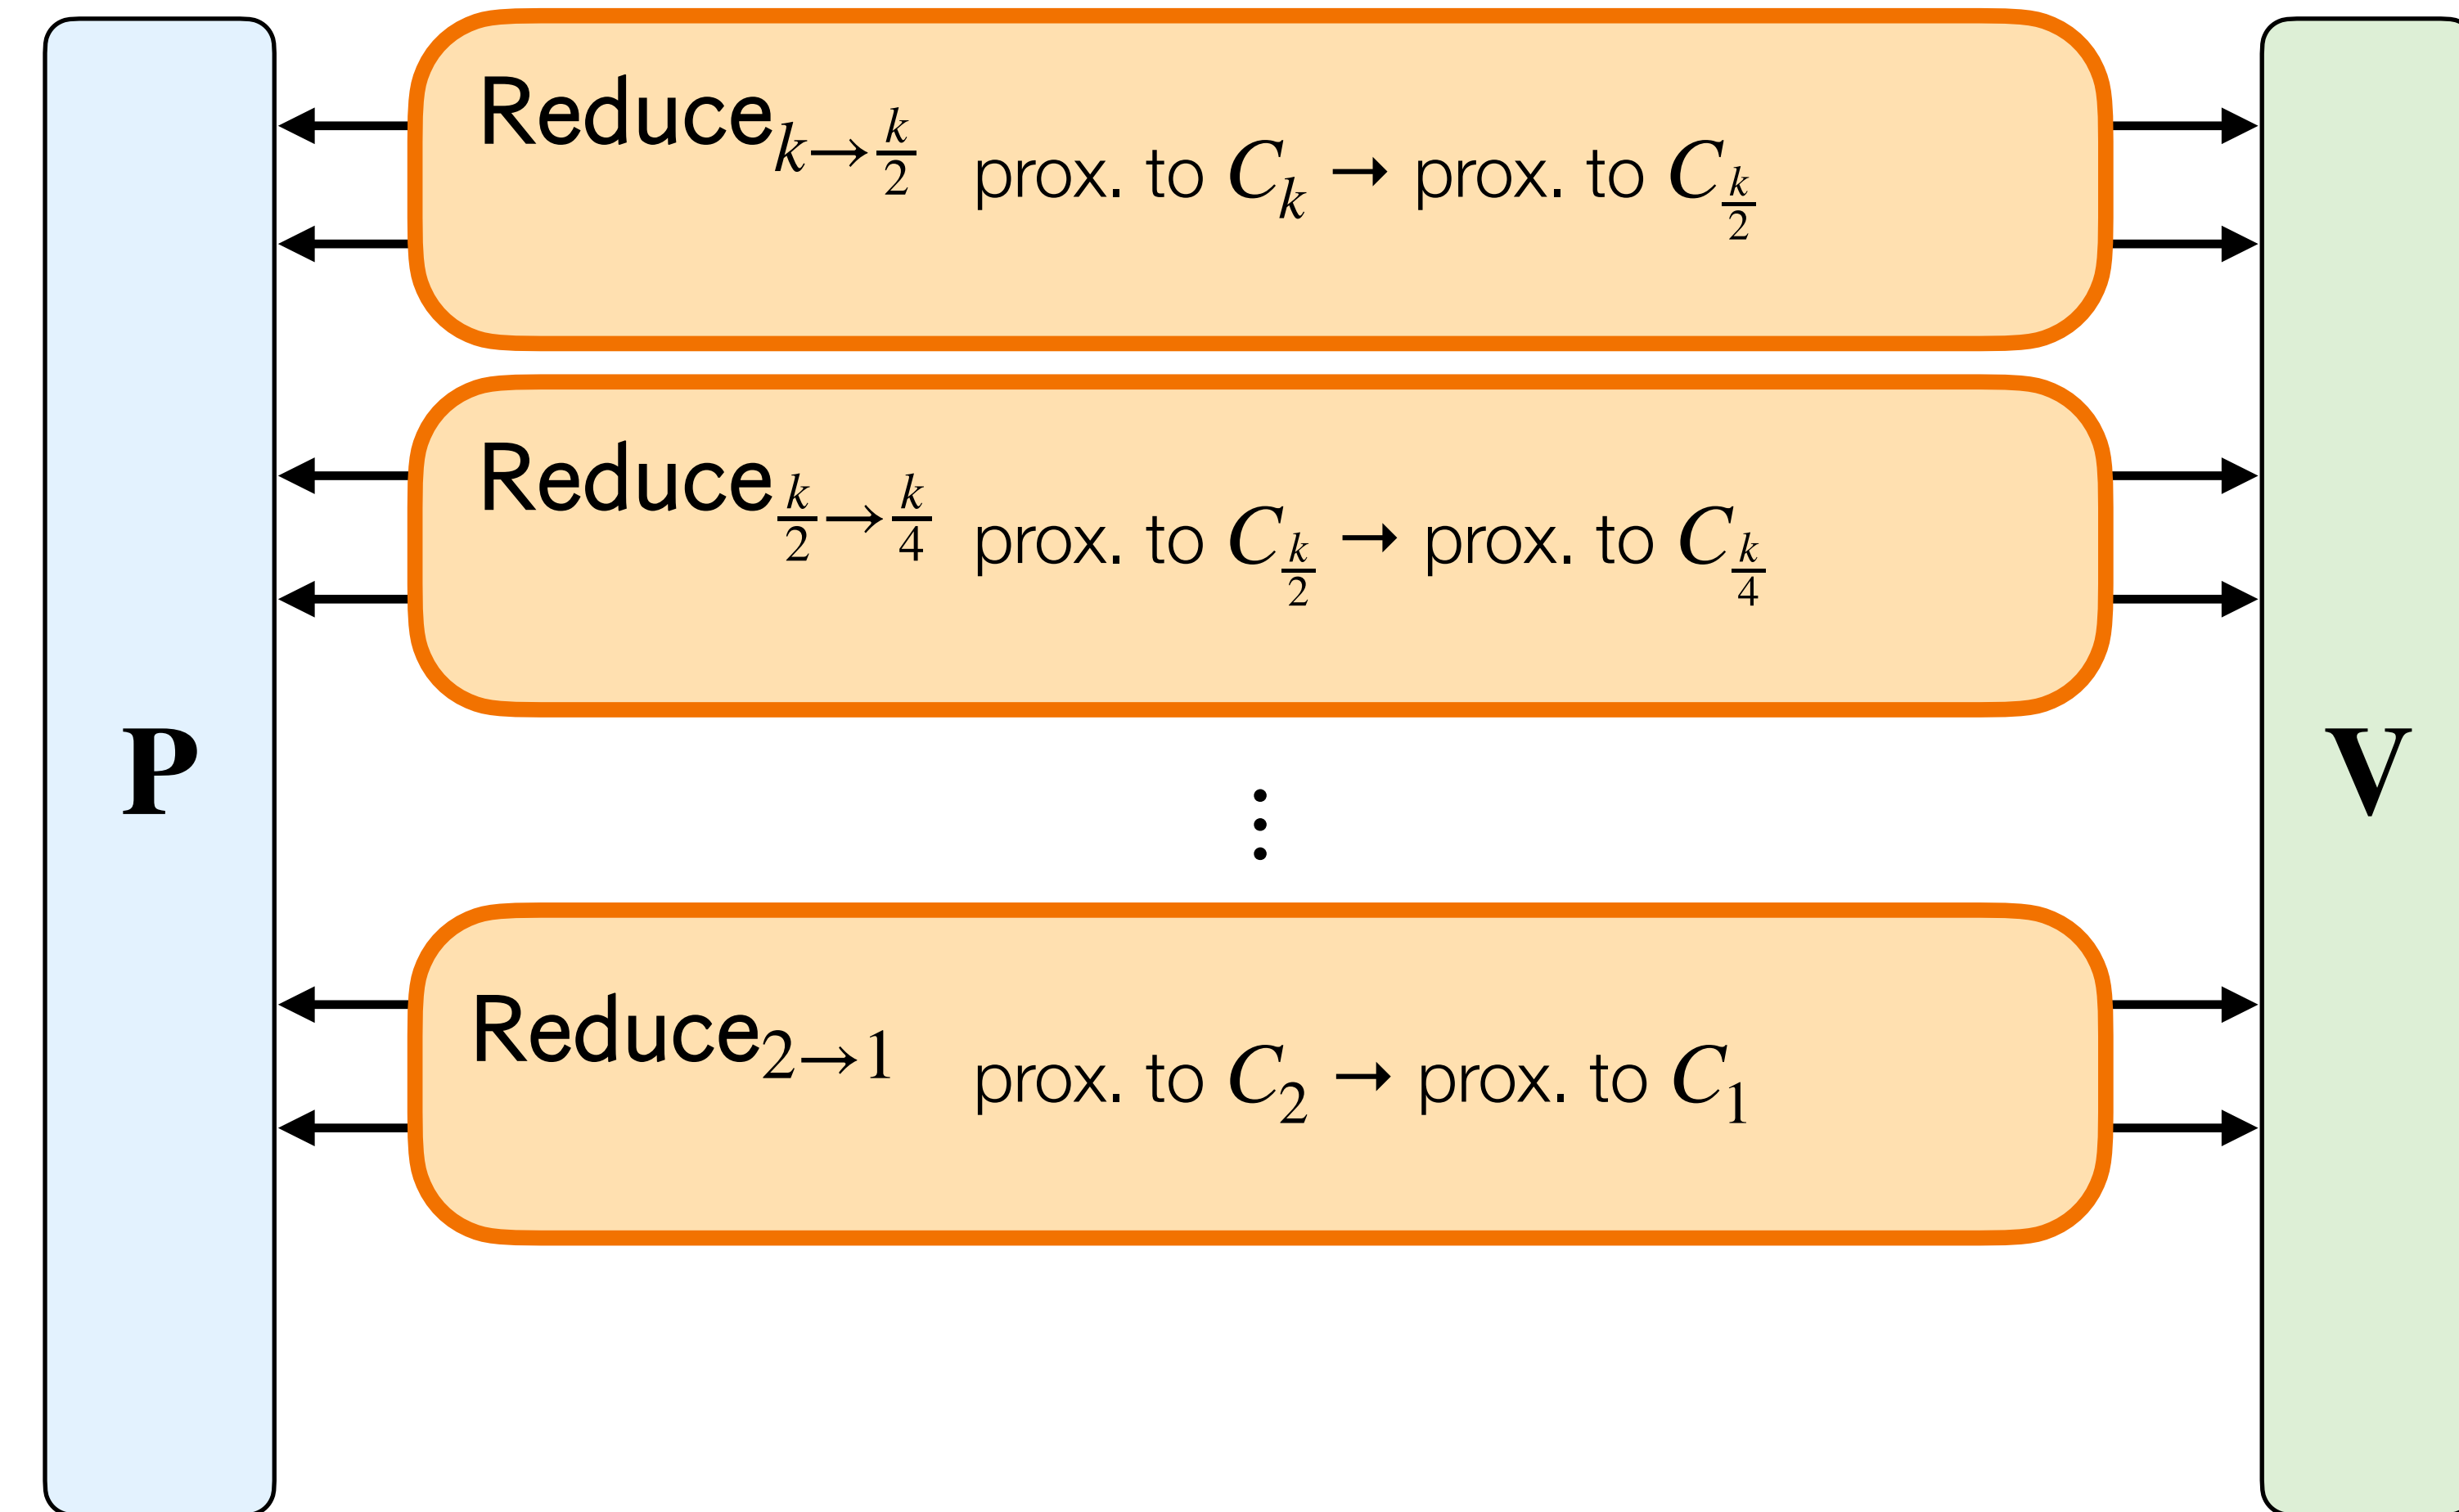
Codeswitch

prox to $B_{\frac{k}{2^{i+1}}} \rightarrow$ prox to $C_{\frac{k}{2^{i+1}}}$

hop to foldable code

$t = 2$

IOPP Completeness and Soundness



Follow from *composition* theorems for interactive oracle reductions

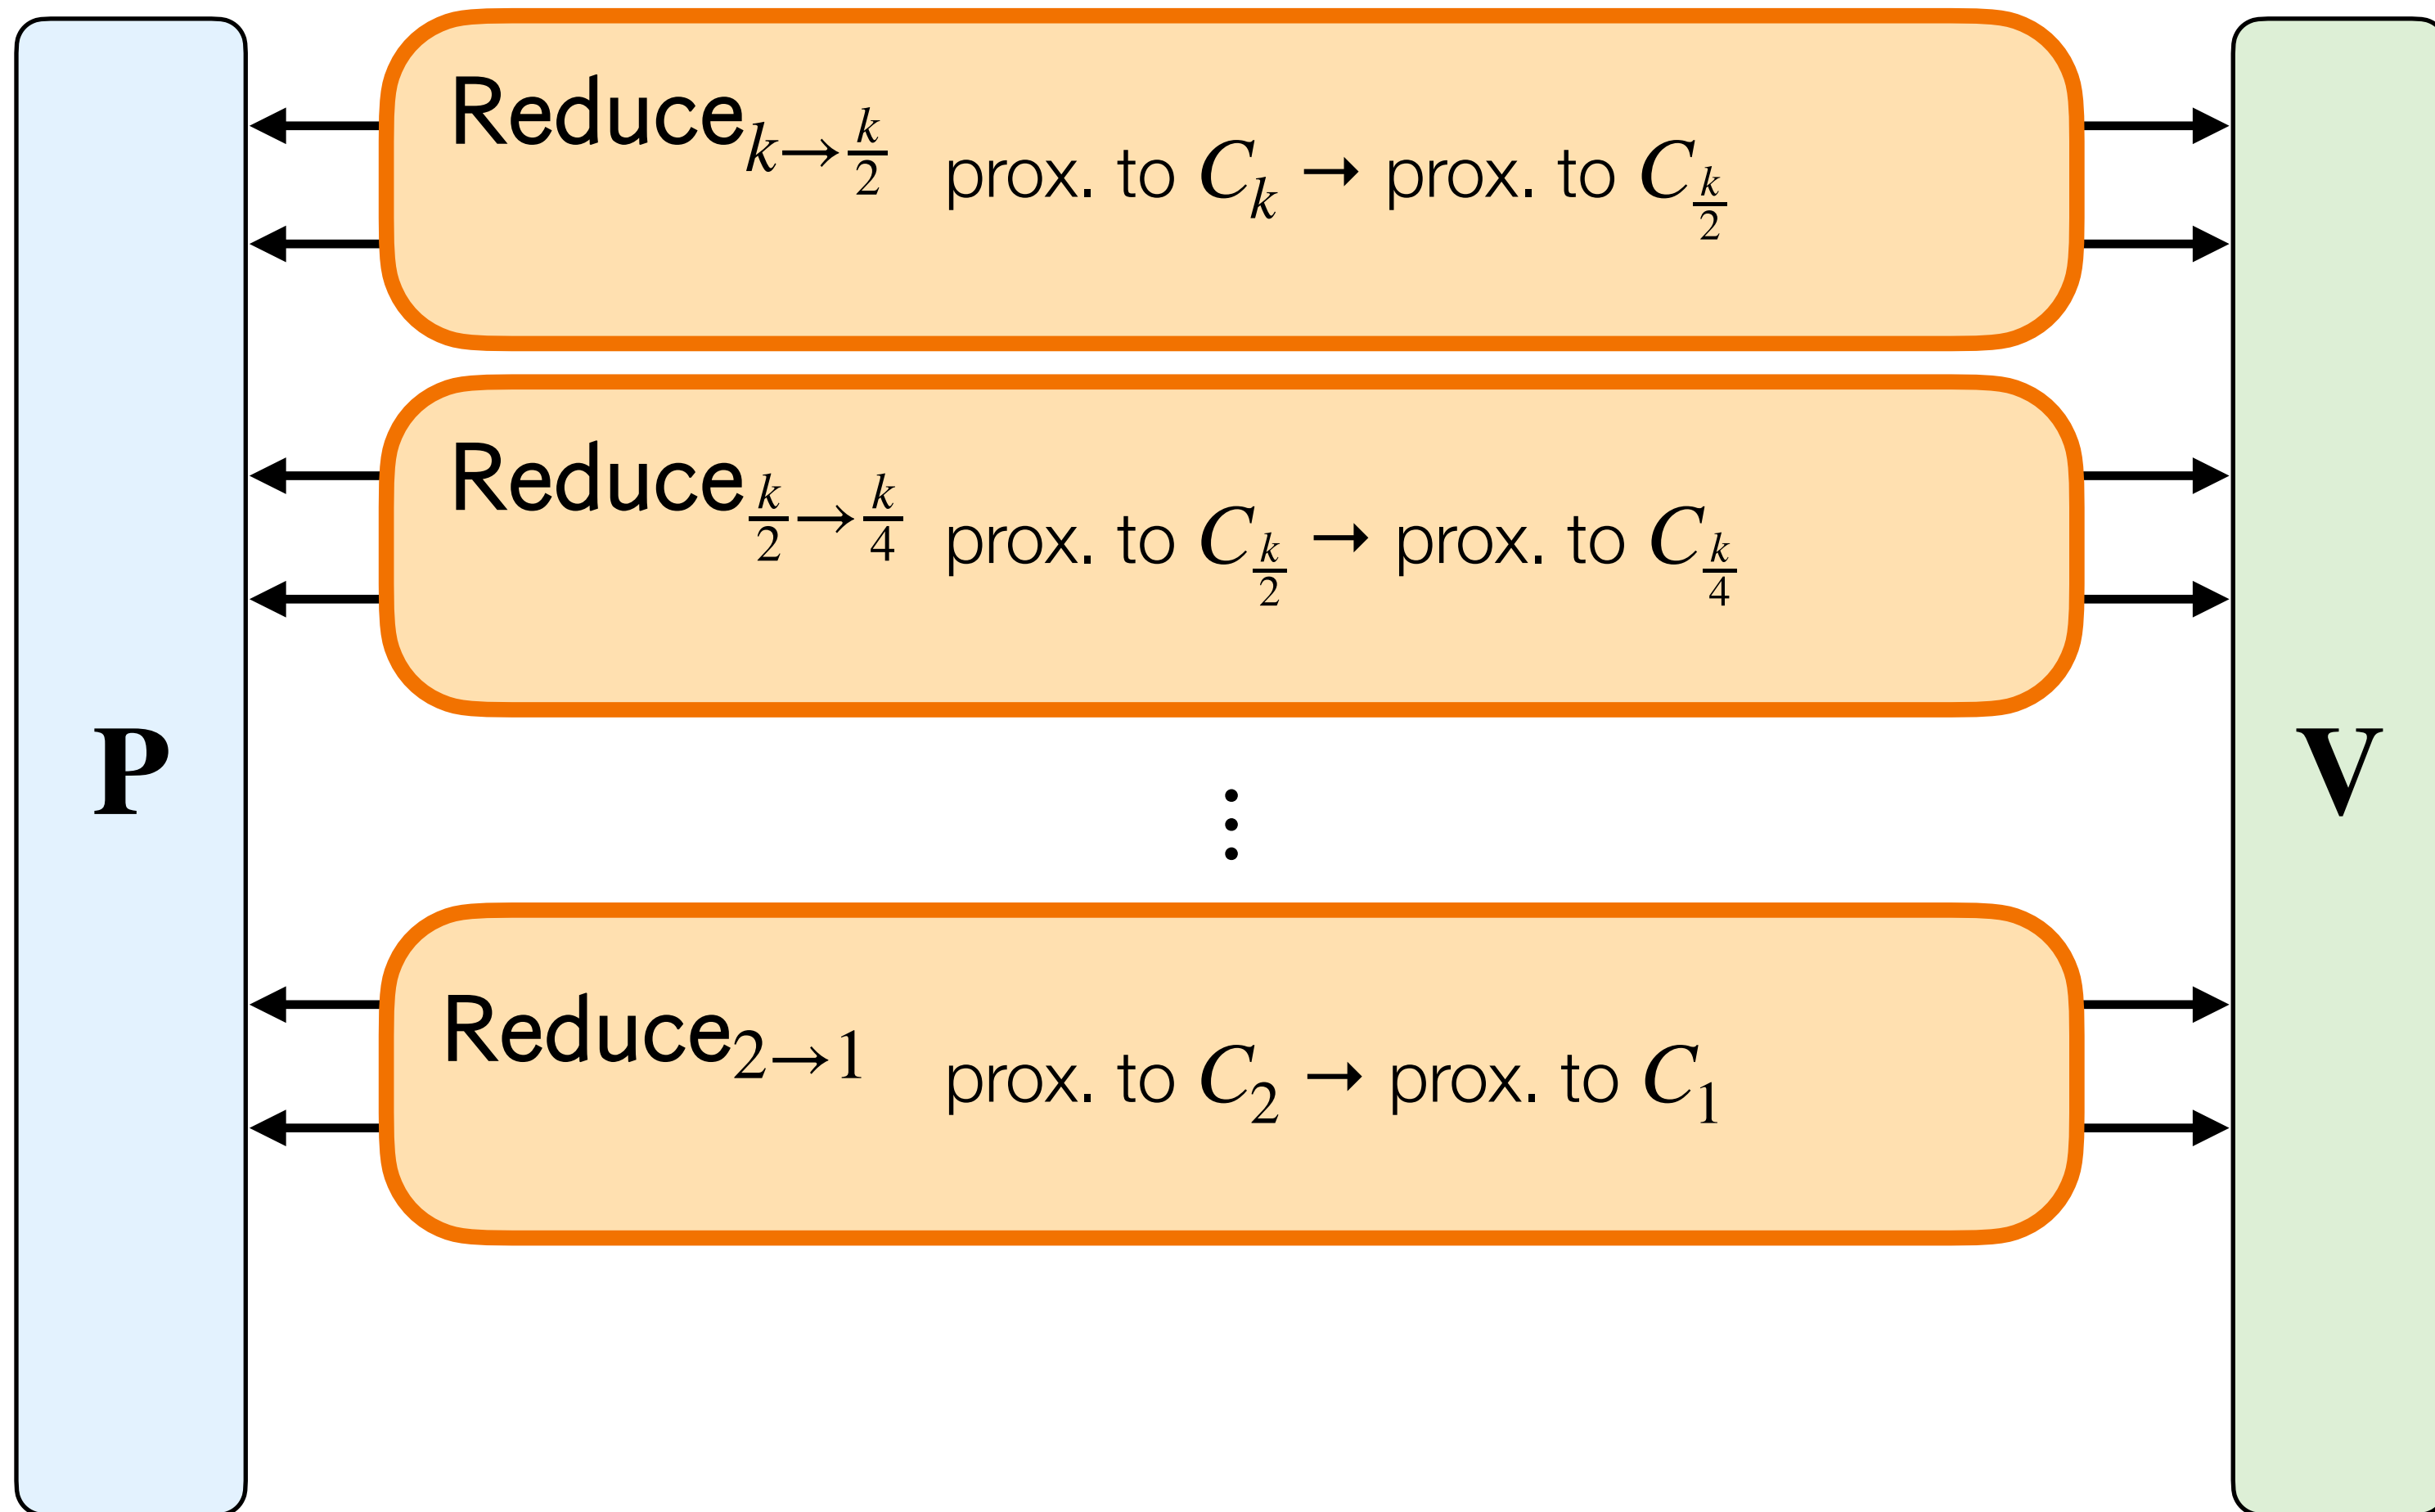
[**BMNW**25, **BMMS**26b]

$t = 2$

IOPP Efficiency

prover time

queries



$$T_{C_k}$$

λ queries

+

+

$$T_{C_{\frac{k}{2}}}$$

λ queries

+

+

⋮

⋮

+

+

$$O(1)$$

λ queries

||

||

$$O(T_{C_k})$$

$$O(\lambda \log k)$$

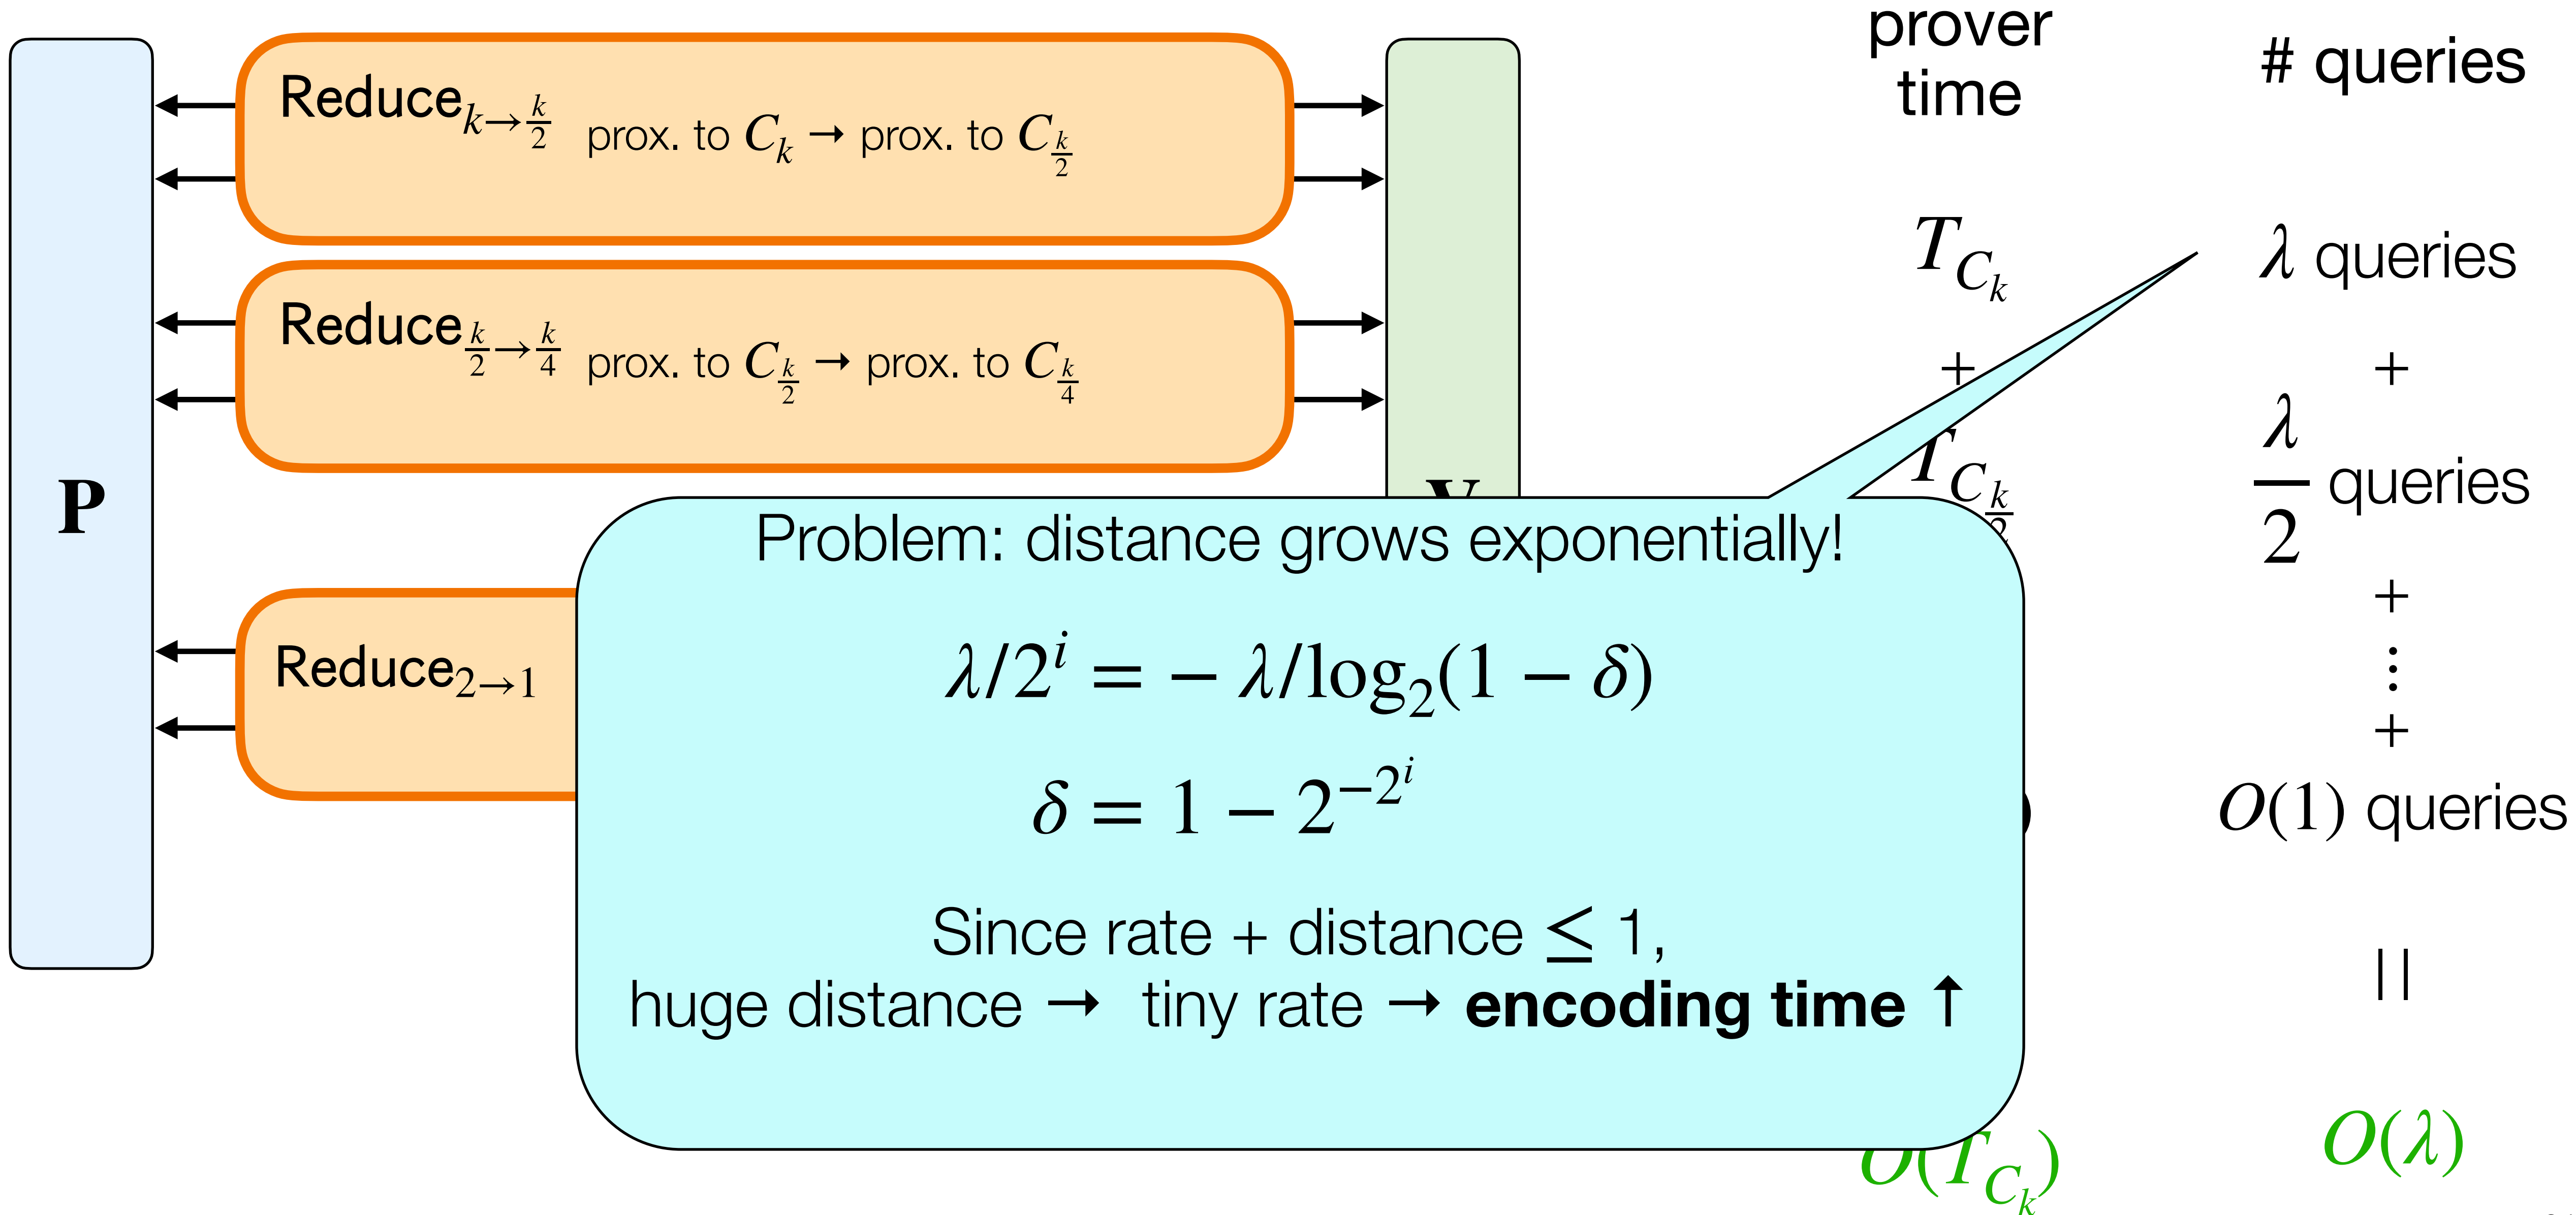
Chapter 3

Query Reduction via Improving Distance


Goal:

IOPP with $O(\lambda \log \log k)$ queries

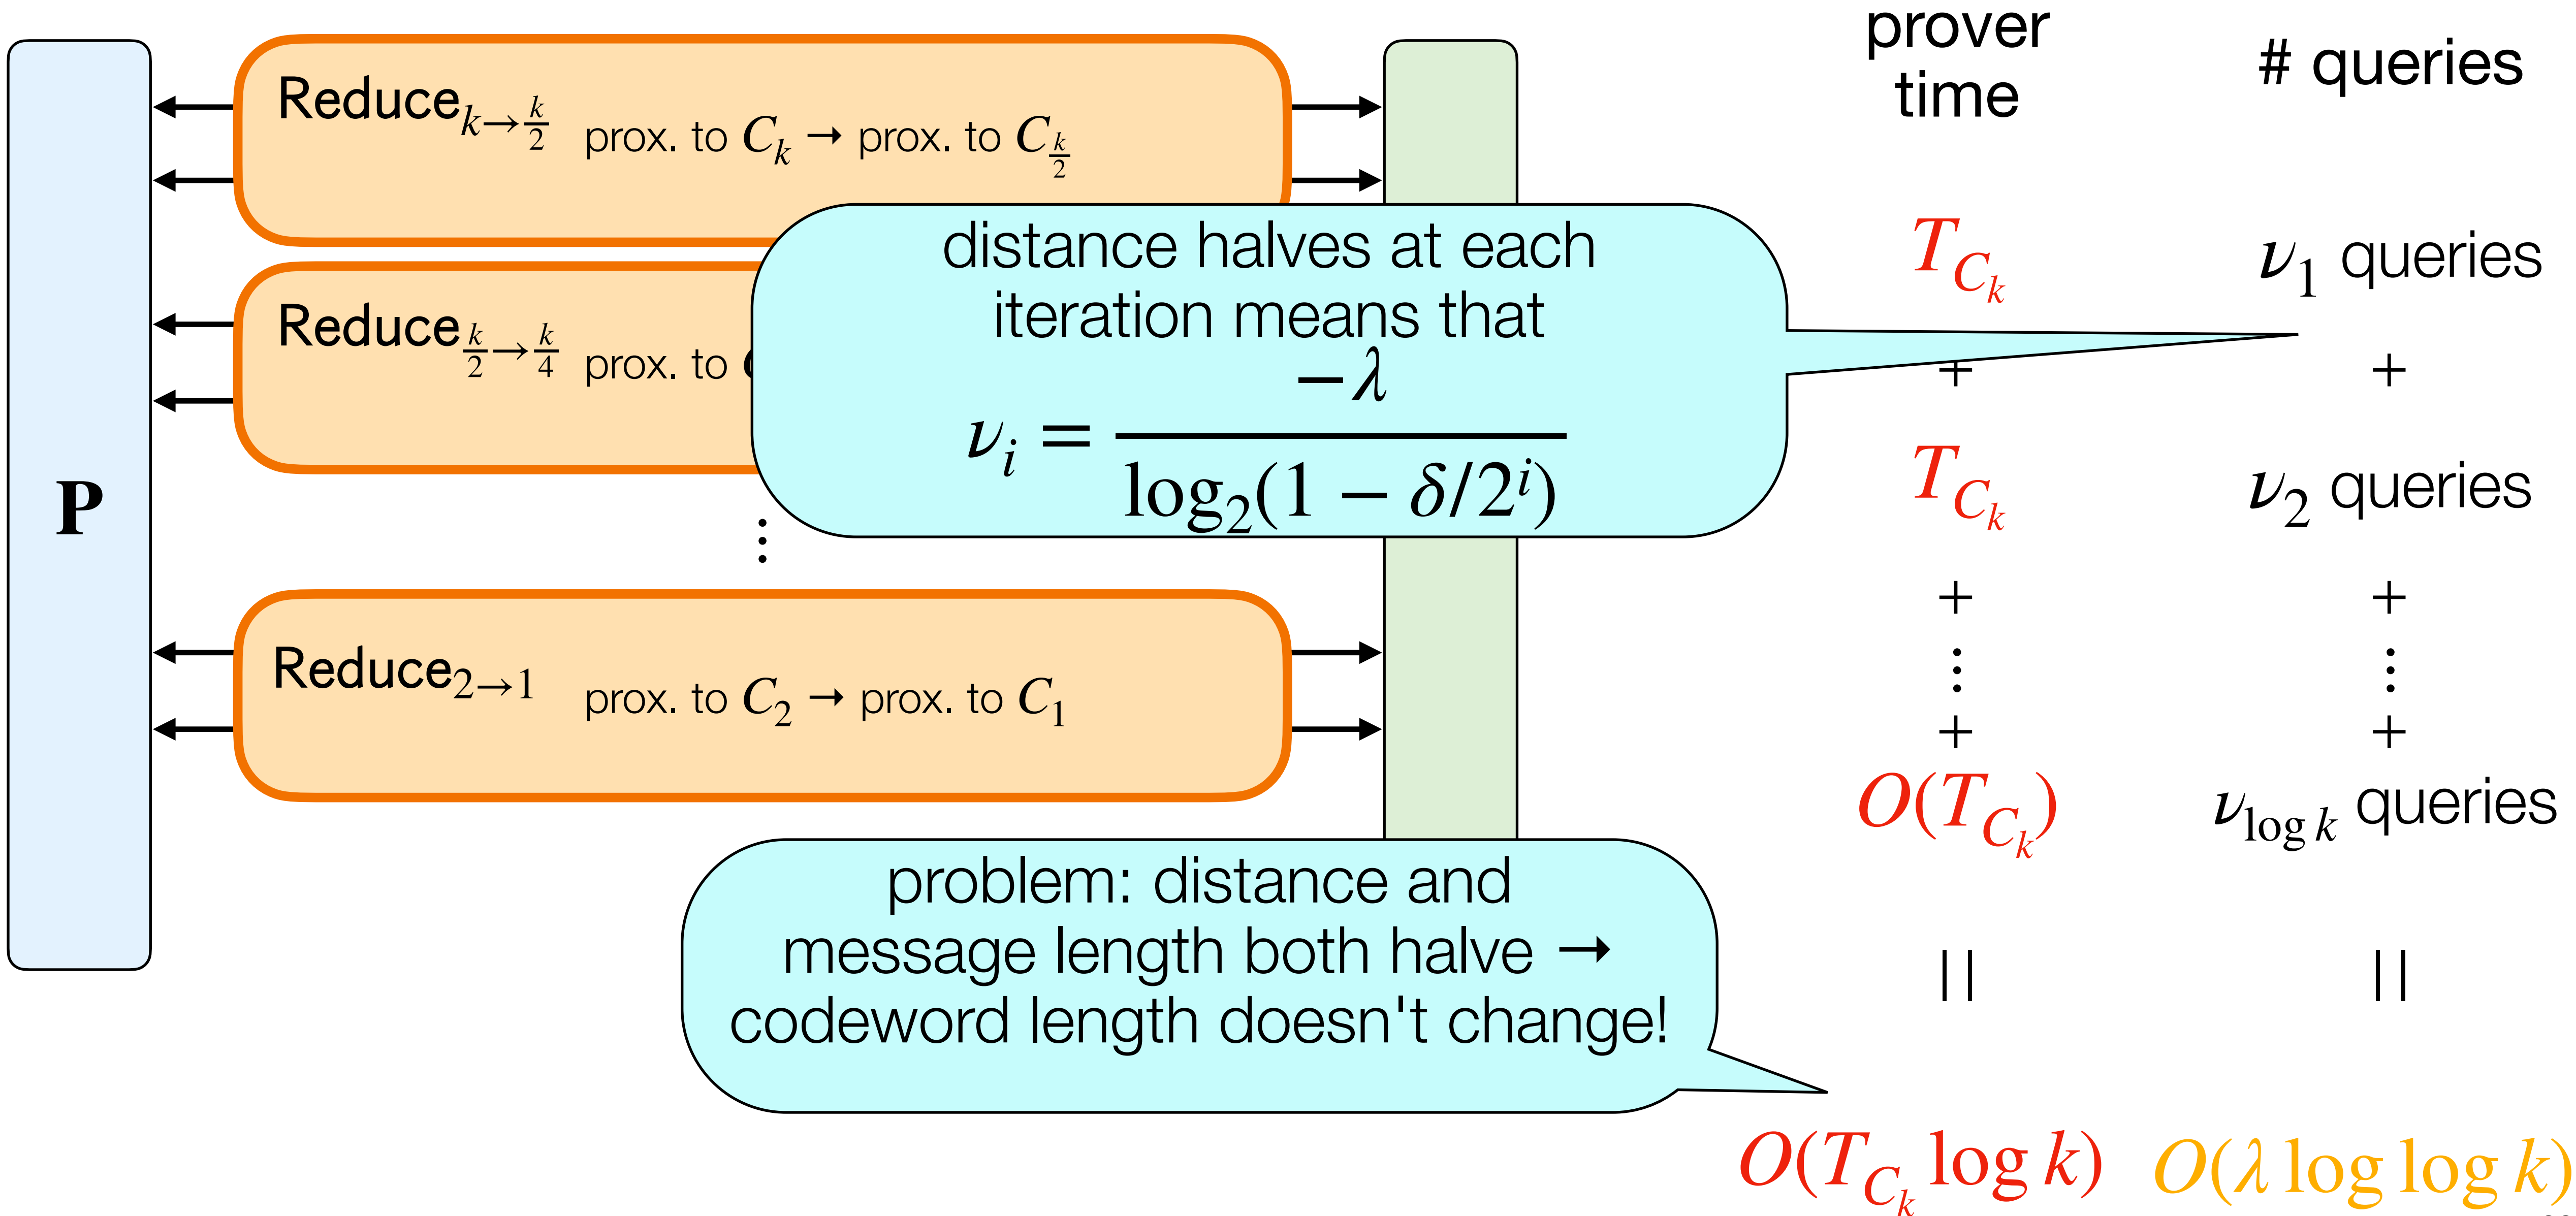
Reduce # queries in each round



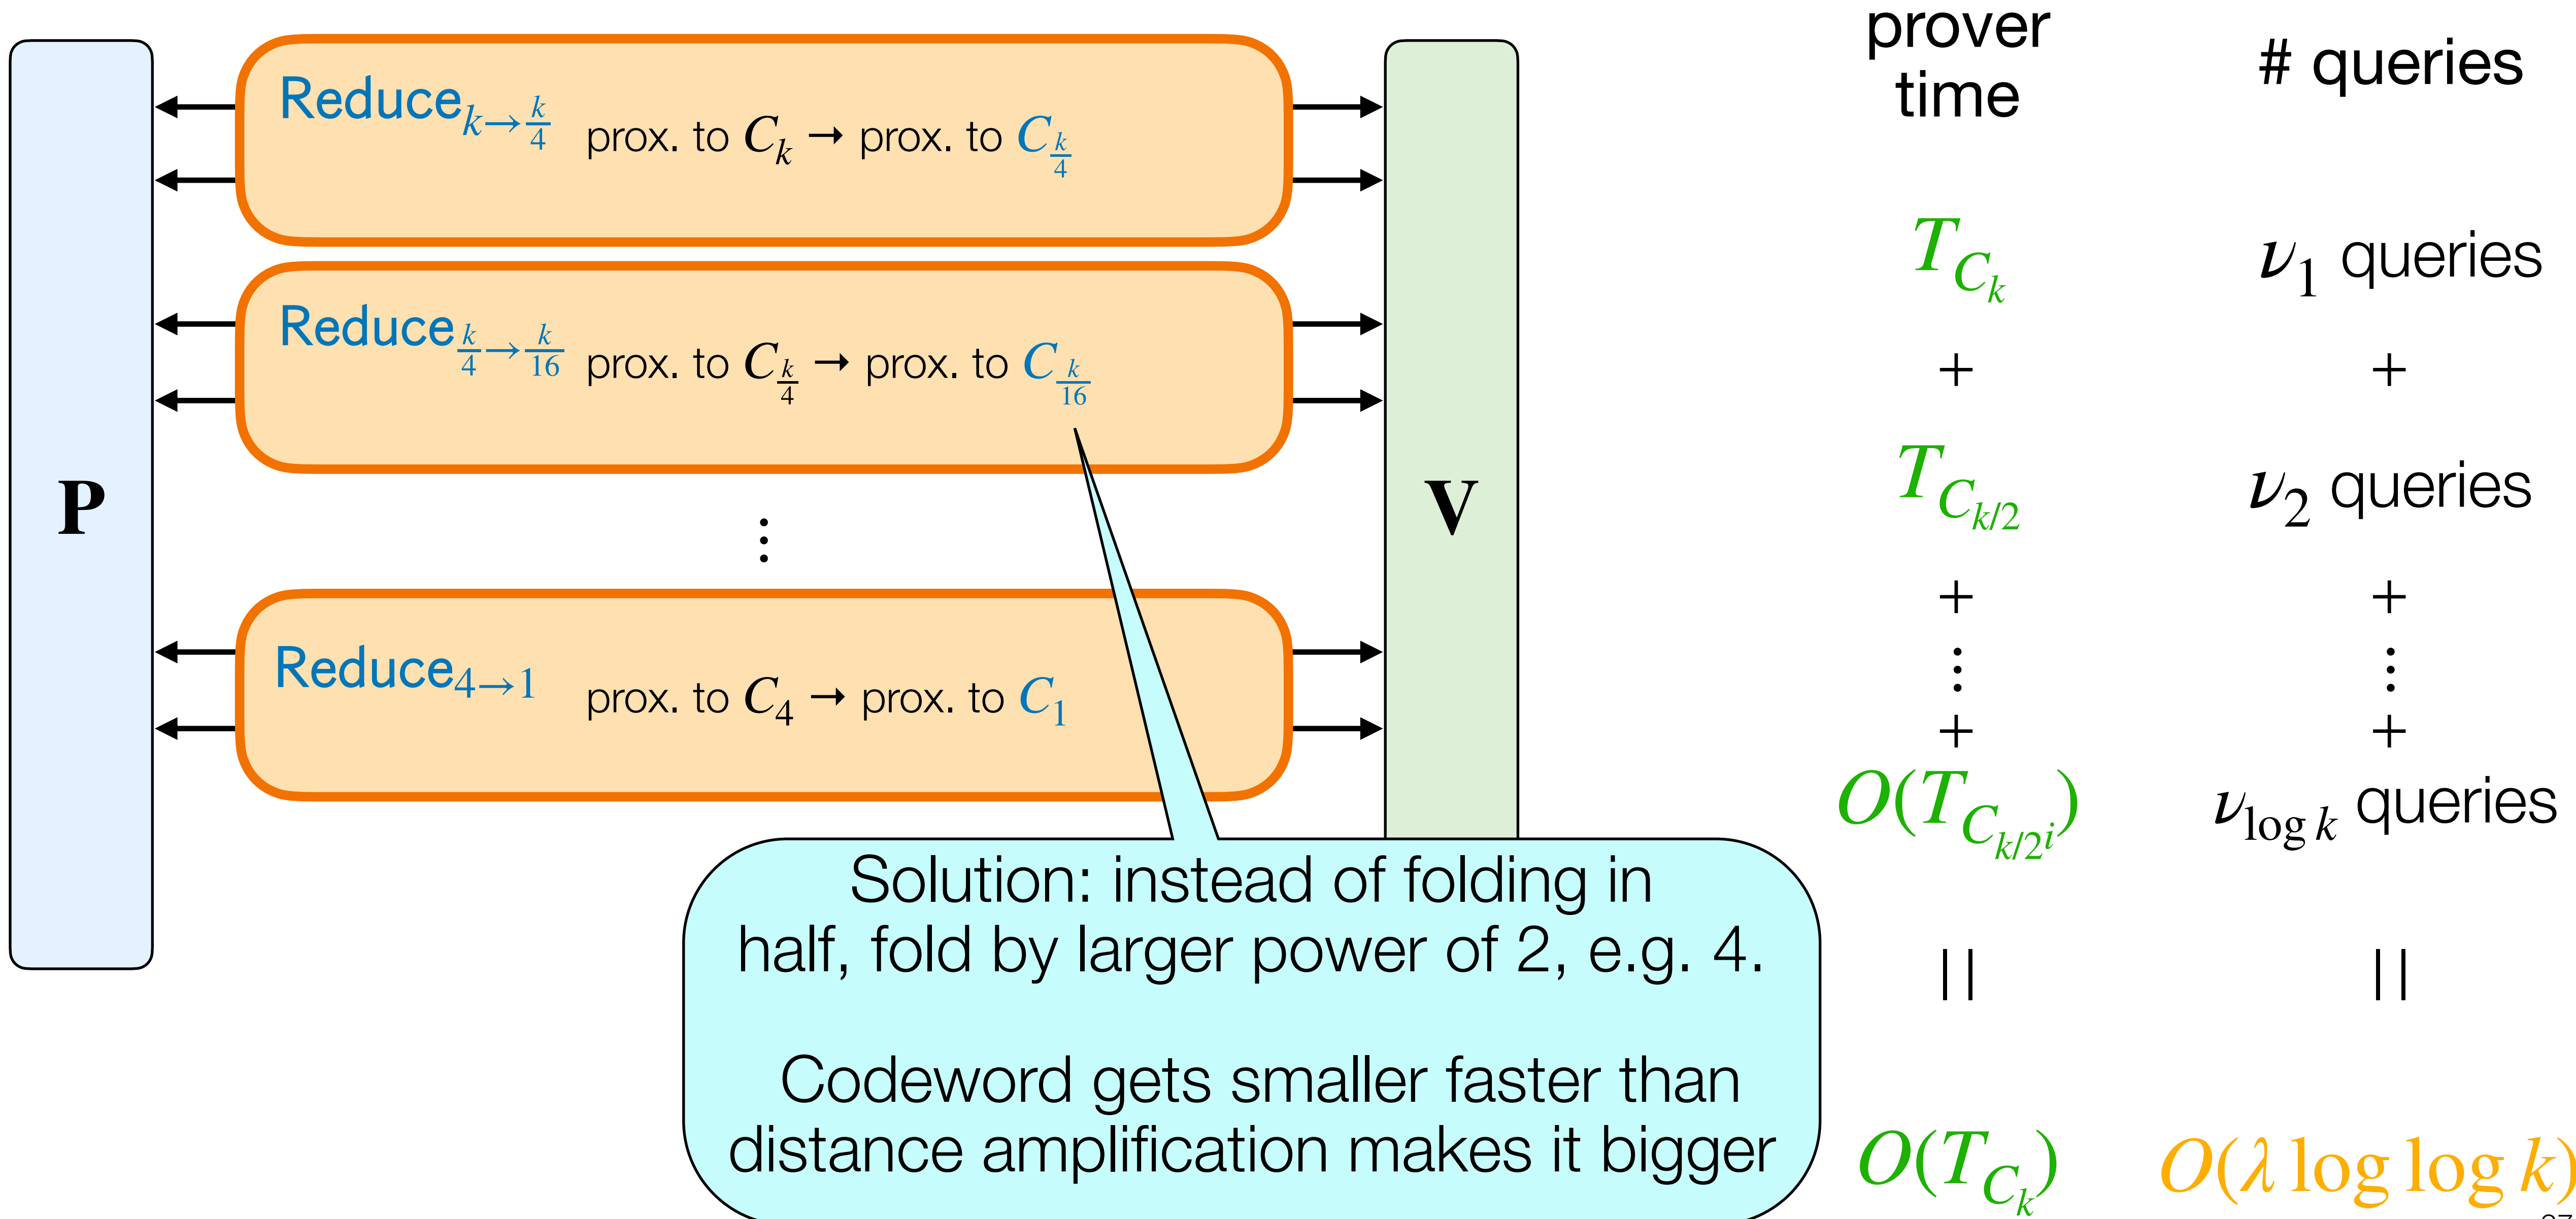
Reduce # queries by increasing distance



Problem: codewords don't shrink!



Solutions: fold more!



Chapter 4

Round *and* Query Reduction via Tensoring

Goal:

IOPP with $O(\lambda)$ queries

Reducing the Number of Reductions

Problem: Too many reductions, and can't reduce the queries per reduction enough!

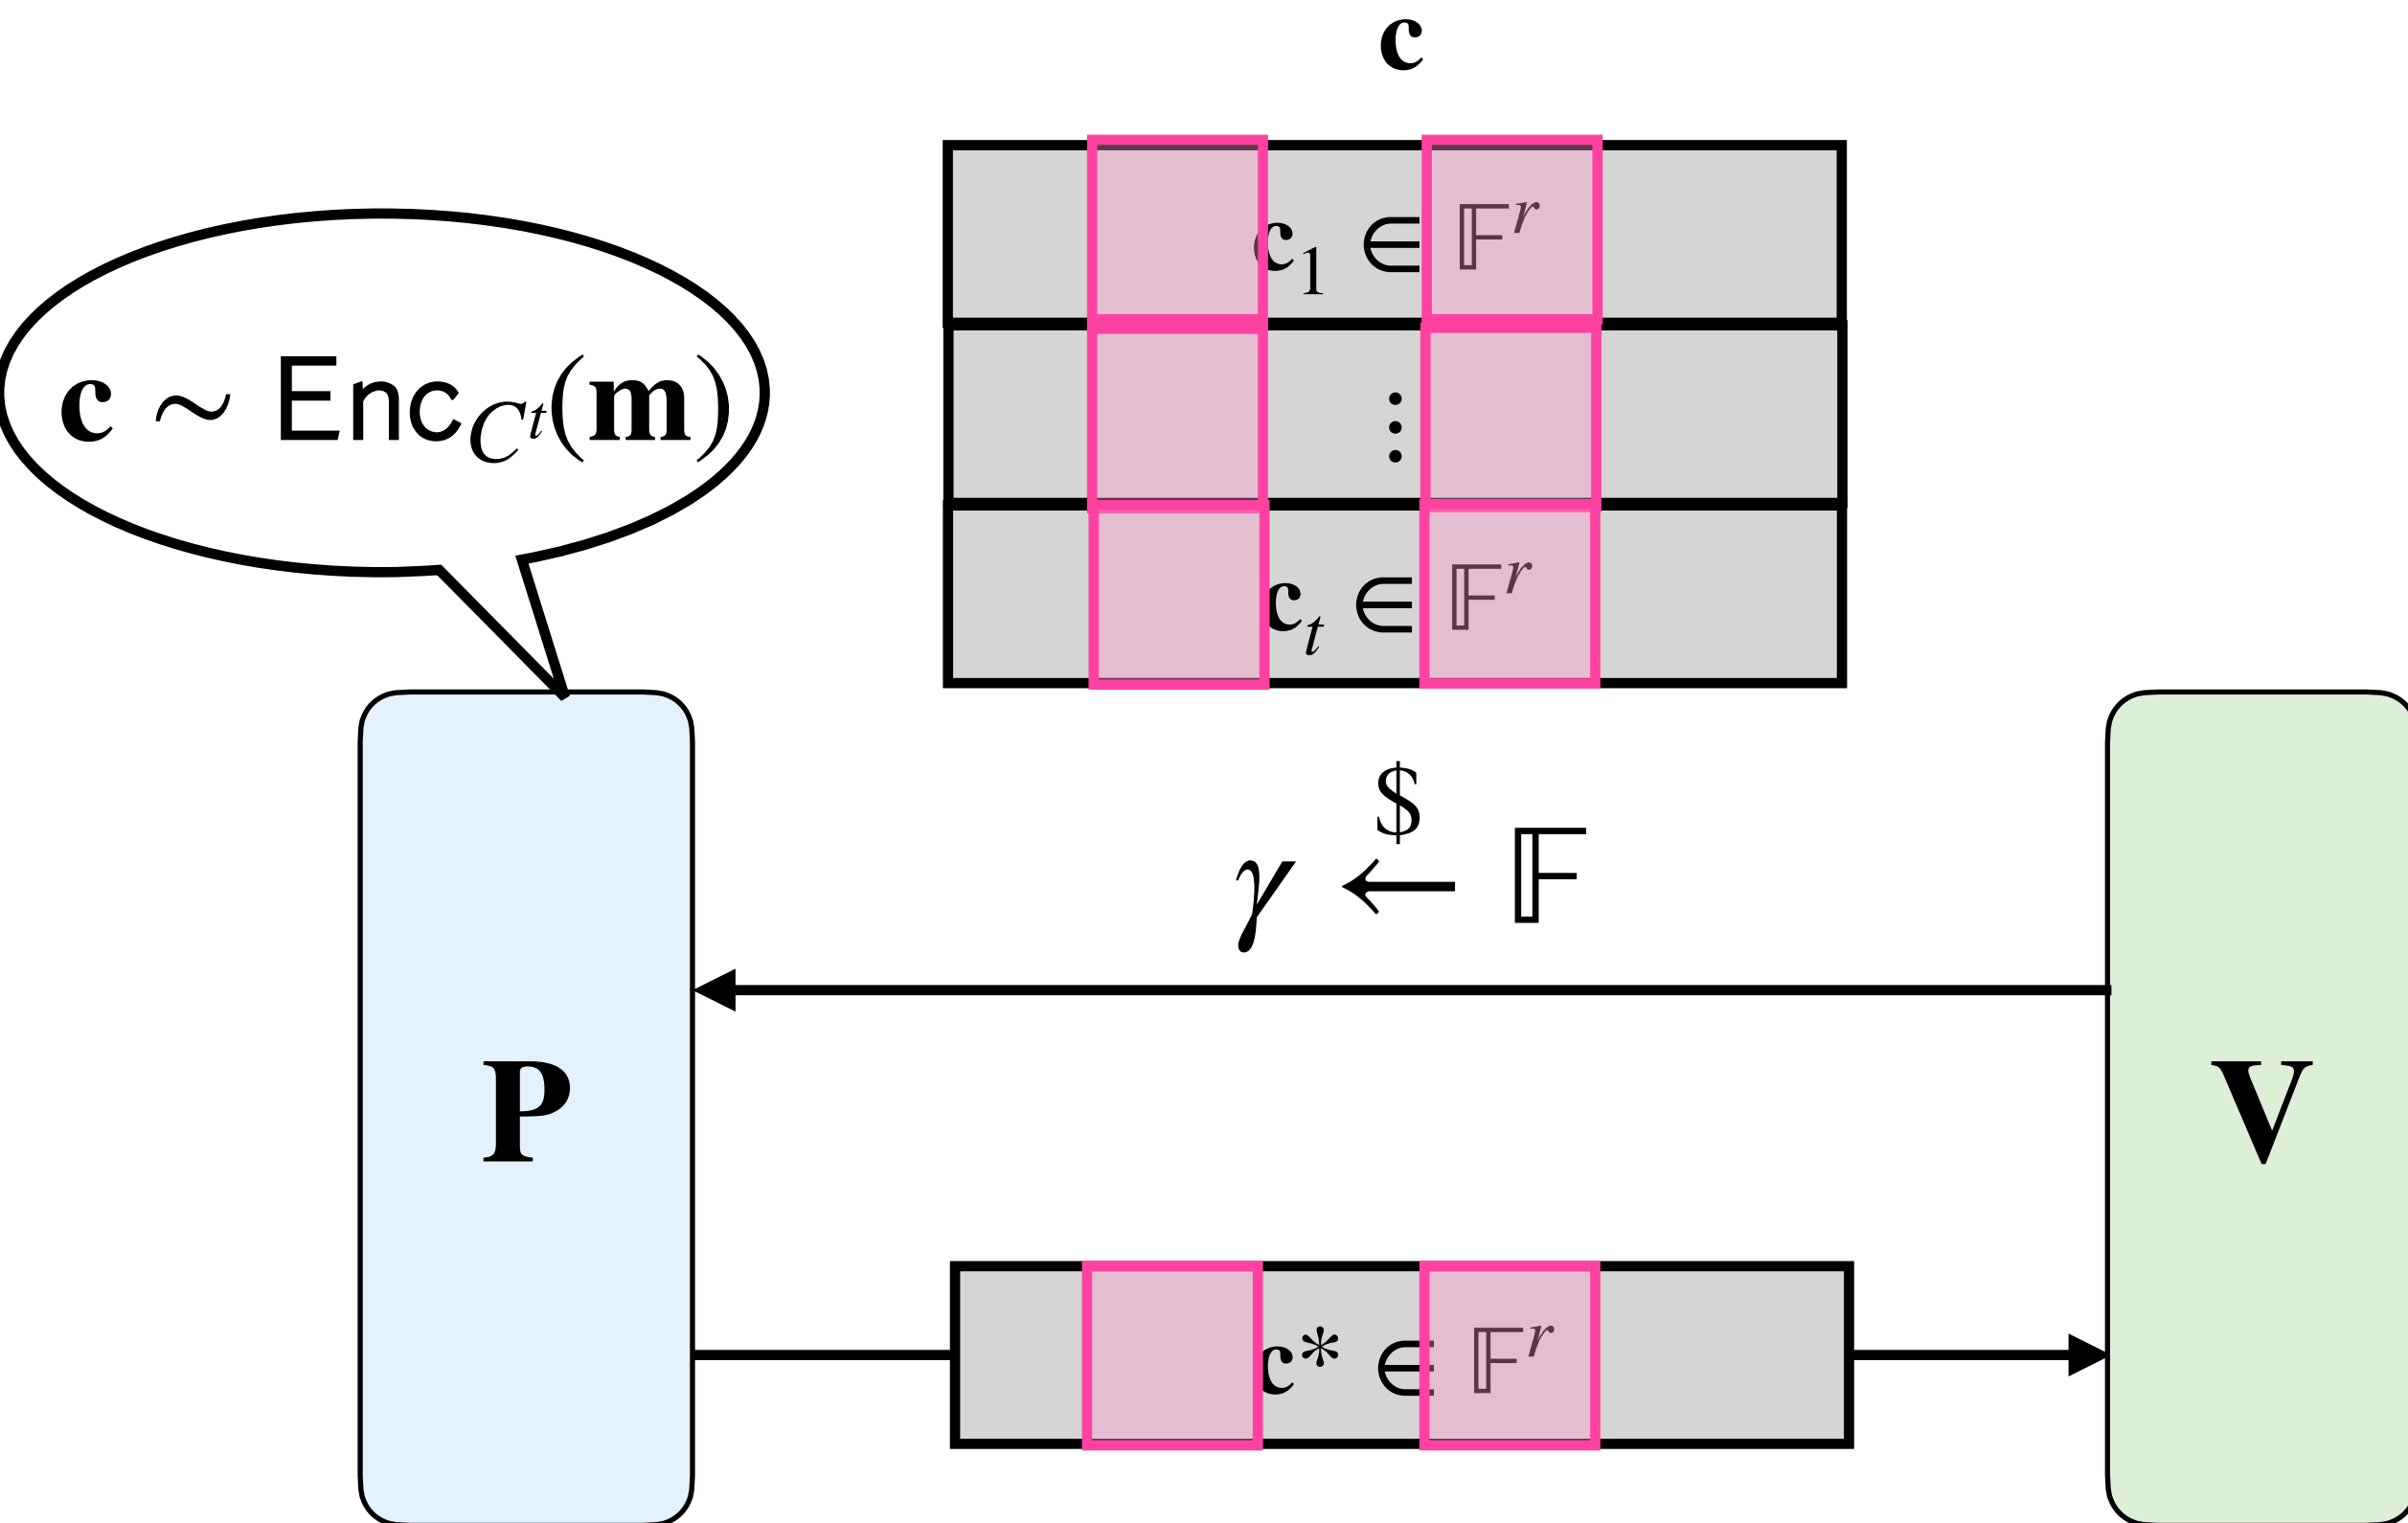
Need: $O(1)$ reductions each with $O(\lambda)$ queries

Good news: suffices to reduce message length to $O(\lambda k^{1/3})$

and codeswitch to RS code with rate $\frac{1}{k^{1/3}}$.

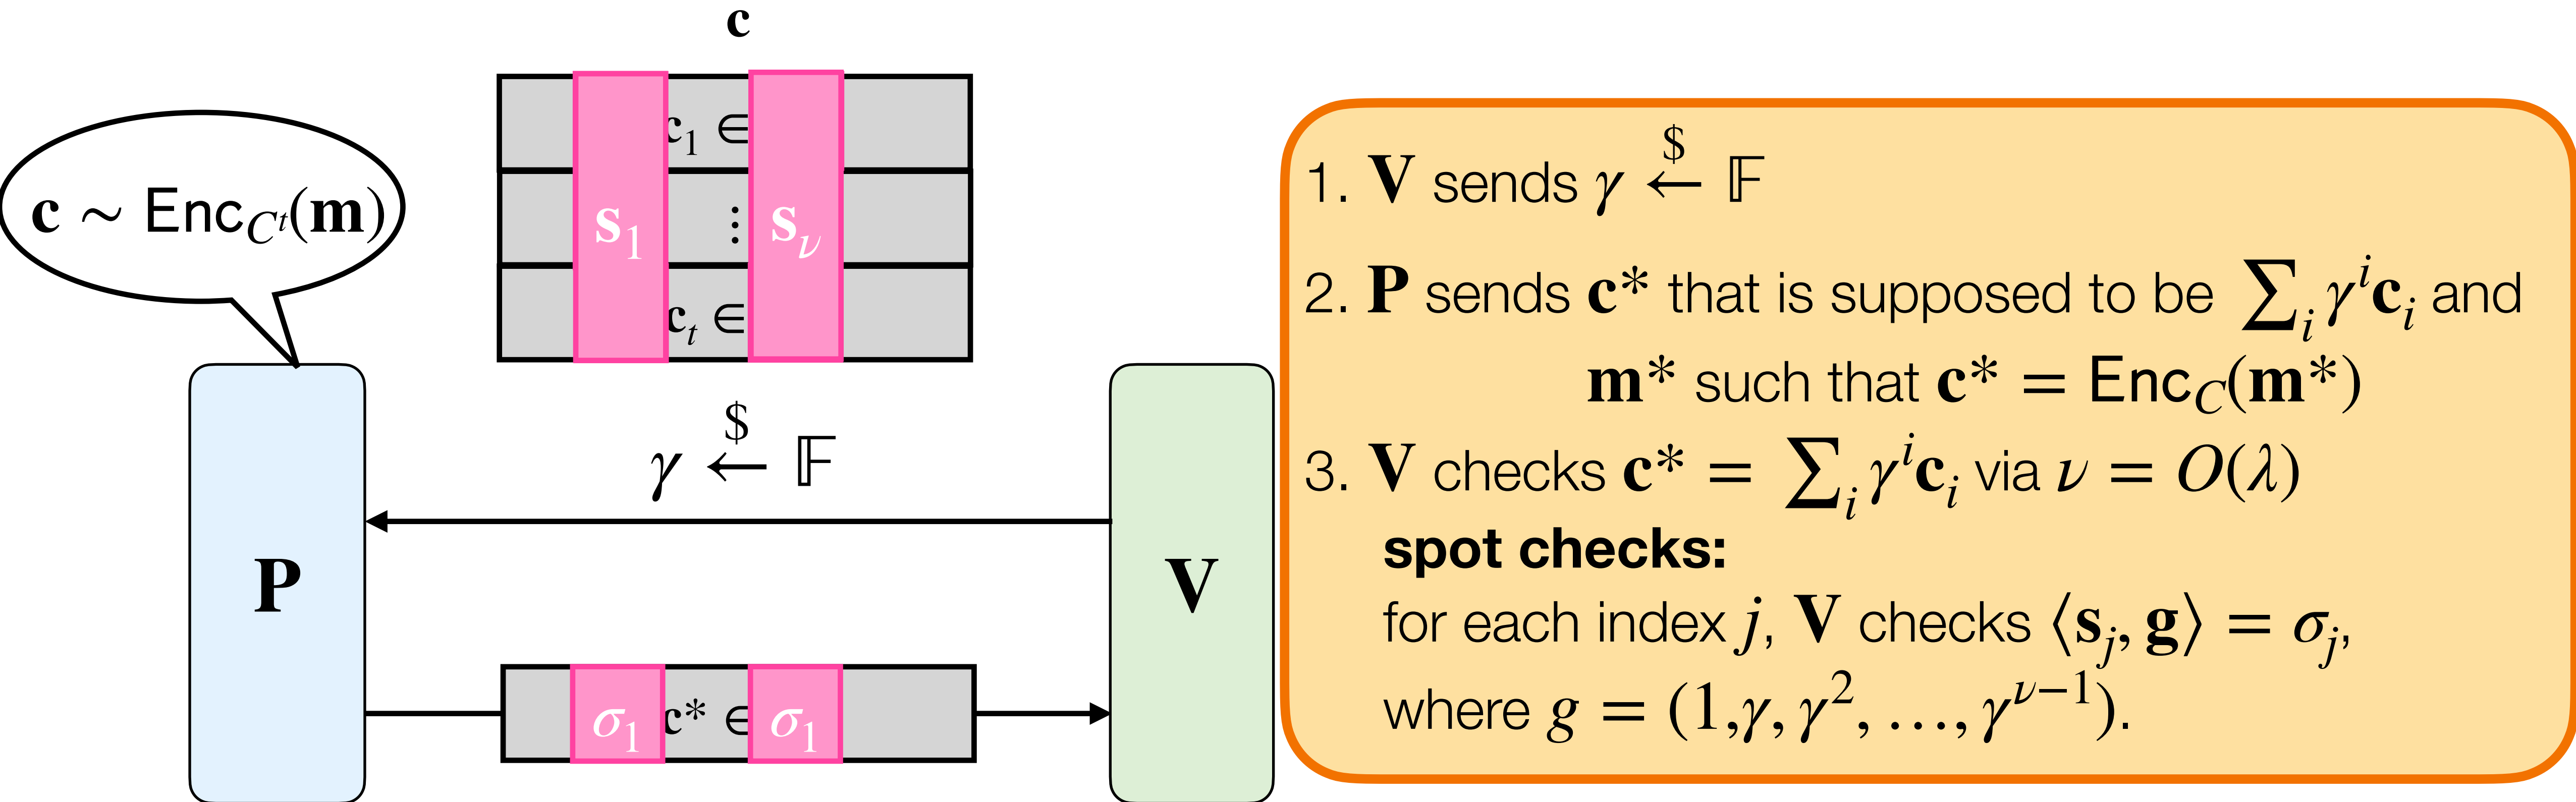
Goal: Reduce $k \rightarrow k^{2/3}$

Idea 1: Interleave with $t = k^{1/3}$



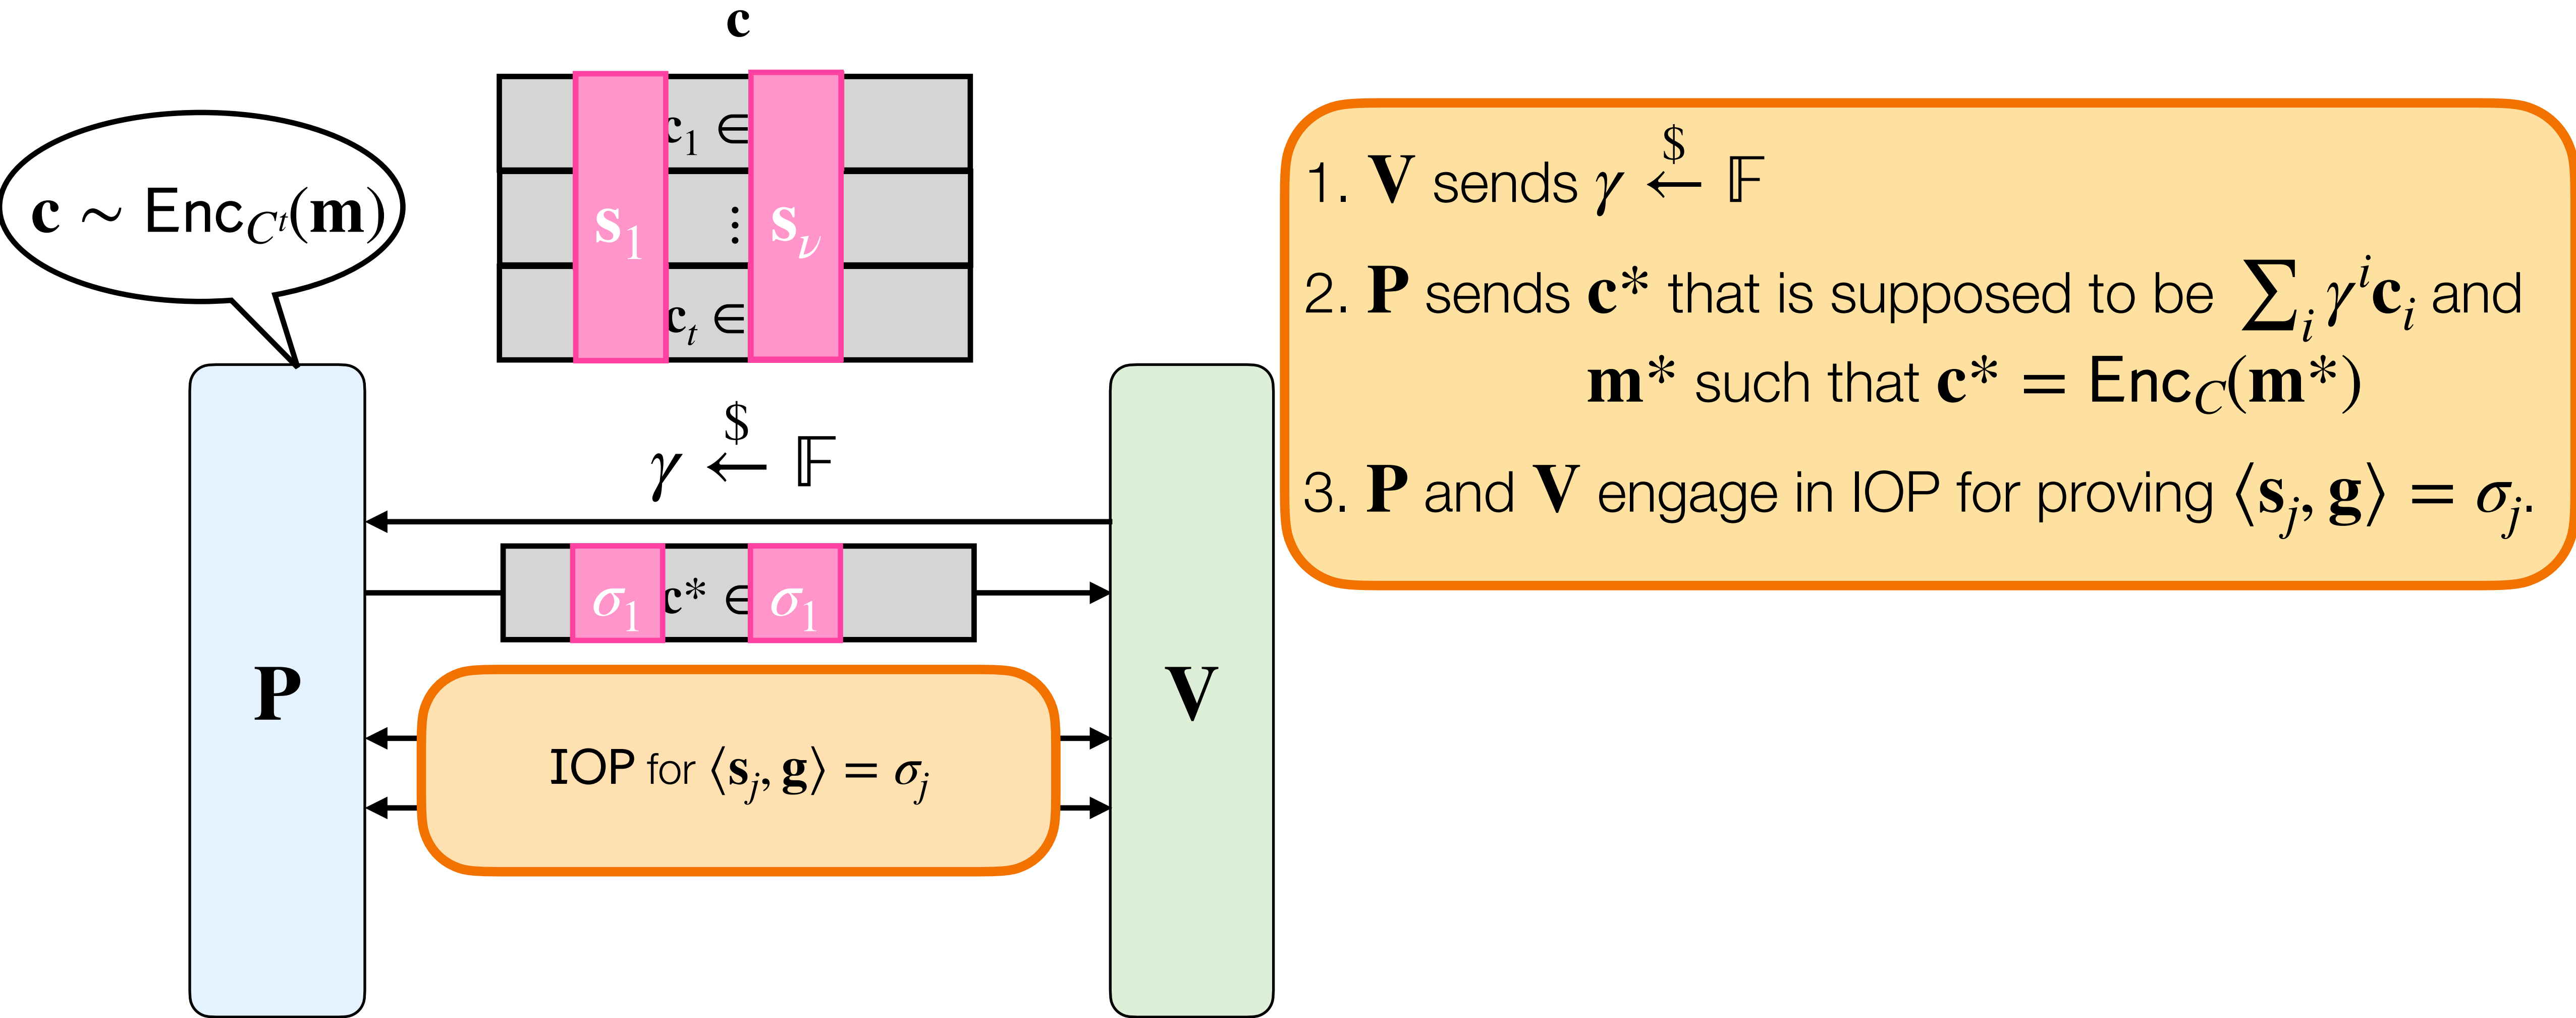
1. \mathbf{V} sends $\gamma \in \mathbb{F}$
2. \mathbf{P} sends \mathbf{c}^* that is supposed to be $\sum_i \gamma^i \mathbf{c}_i$ and \mathbf{m}^* such that $\mathbf{c}^* = \text{Enc}_C(\mathbf{m}^*)$
3. \mathbf{V} checks $\mathbf{c}^* = \sum_i \gamma^i \mathbf{c}_i$ via $\nu = O(\lambda)$
spot checks:
 1. \mathbf{V} samples ν indices j_1, \dots, j_ν .
 2. For each index j :
 1. \mathbf{V} queries $\mathbf{c}_1[j], \dots, \mathbf{c}_t[j]$ and $\mathbf{c}^*[j]$
 2. \mathbf{V} checks that $\mathbf{c}^*[j] = \sum_i \gamma^i \cdot \mathbf{c}_i[j]$

Idea 1: Interleave with $t = k^{1/3}$



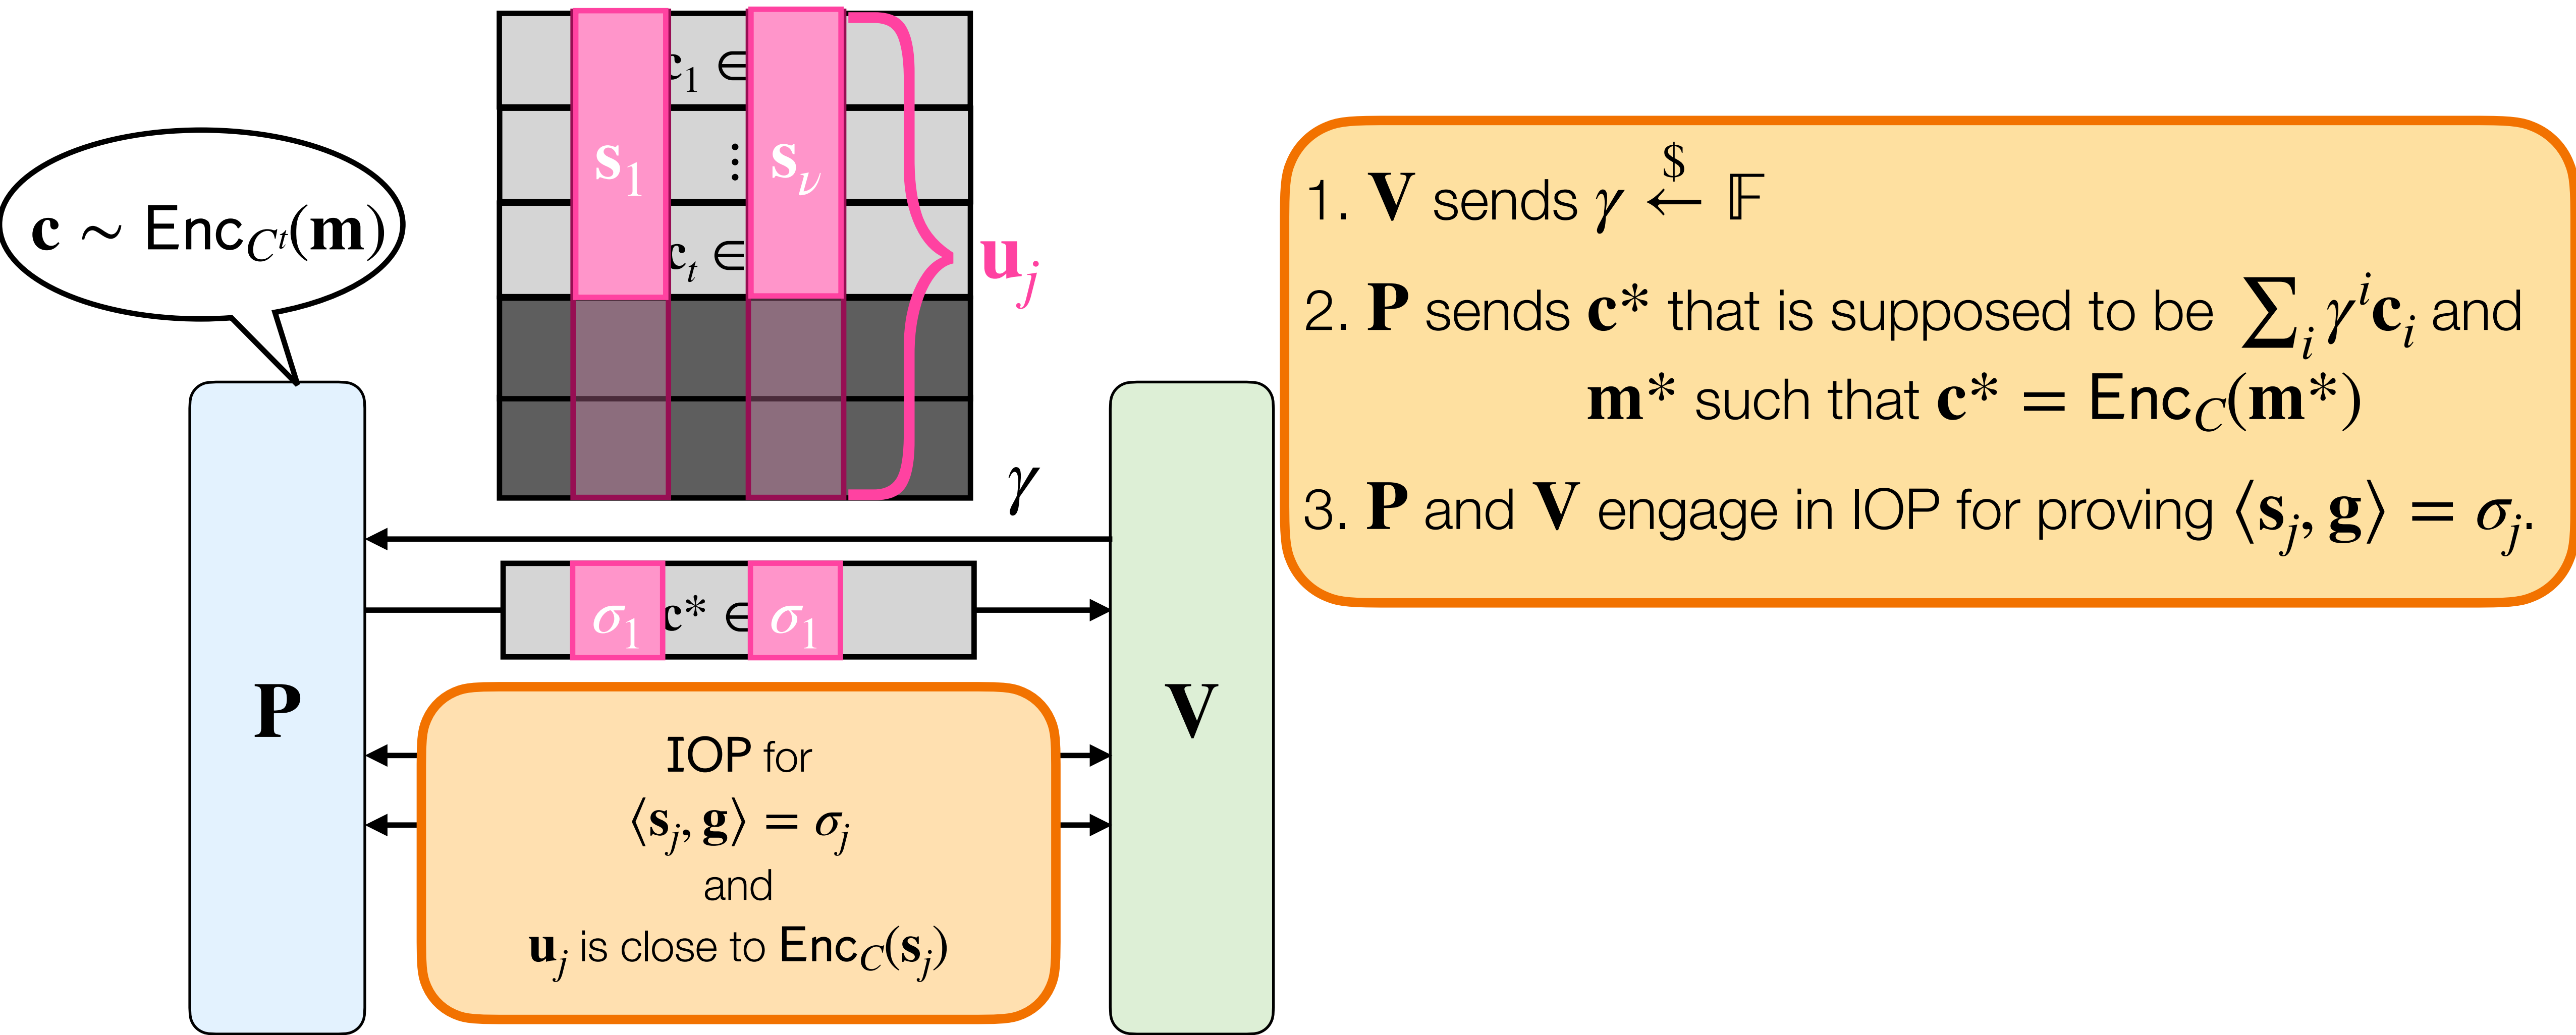
Problem: $O(\lambda \cdot k^{1/3})$ queries!

Idea 2: Delegate spot checks to prover



Problem: **V** must still read input \mathbf{s}_j 's


Idea 2.1: Encode Columns Too



1. \mathbf{V} sends $\gamma \xleftarrow{\$} \mathbb{F}$
2. \mathbf{P} sends \mathbf{c}^* that is supposed to be $\sum_i \gamma^i \mathbf{c}_i$ and \mathbf{m}^* such that $\mathbf{c}^* = \text{Enc}_C(\mathbf{m}^*)$
3. \mathbf{P} and \mathbf{V} engage in IOP for proving $\langle \mathbf{s}_j, \mathbf{g} \rangle = \sigma_j$.


Problem: Naive IOP still incurs $O(\lambda^2)$ queries

IOP for Inner-Products and Proximity



1. \mathbf{P} sends $\mathbf{u} = \text{Enc}_D(\mathbf{s}_1 \parallel \dots \parallel \mathbf{s}_\nu)$ under code D
2. \mathbf{P} convinces \mathbf{V} that:
 1. \mathbf{u} is actually an encoding of $\mathbf{s}_1 \parallel \dots \parallel \mathbf{s}_\nu$:
 1. \mathbf{u} is close to D
 2. Message in \mathbf{u} is consistent with \mathbf{u}_j 's
 2. For each j , $\langle \mathbf{s}_j, \mathbf{g} \rangle = \sigma_j$

IOP for Inner-Products and Proximity




1. \mathbf{P} sends $\mathbf{u} = \text{Enc}_{\text{RS}}(\mathbf{s}_1 \parallel \dots \parallel \mathbf{s}_\nu)$ under RS code
2. \mathbf{P} convinces \mathbf{V} that:
 1. \mathbf{u} is actually an encoding of $\mathbf{s}_1 \parallel \dots \parallel \mathbf{s}_\nu$:
 1. \mathbf{u} is close to D
 2. Message in \mathbf{u} is consistent with \mathbf{u}_j 's
 2. For each j , $\langle \mathbf{s}_j, \mathbf{g} \rangle = \sigma_j$

RS code of msg length $\nu k^{1/3}$ & block length $k^{1/2}$
 (this is small enough that encoding is still $\ll k$).

WHIR proves proximity and sumcheck claims; in
 this distance regime, WHIR requires $O(\lambda)$ queries.

IOP for Inner-Products and Proximity




1. \mathbf{P} sends $\mathbf{u} = \text{Enc}_{\text{RS}}(\mathbf{s}_1 \parallel \dots \parallel \mathbf{s}_\nu)$ under RS code
2. \mathbf{P} convinces \mathbf{V} that:
 1. \mathbf{u} is actually an encoding of $\mathbf{s}_1 \parallel \dots \parallel \mathbf{s}_\nu$:
 1. \mathbf{u} is close to D
 2. Message in \mathbf{u} is consistent with \mathbf{u}_j 's
 2. For each j , $\langle \mathbf{s}_j, \mathbf{g} \rangle = \sigma_j$

Attempt 1:

1. \mathbf{V} samples λ random indices per \mathbf{u}_j .
2. For each index i , \mathbf{P} proves that $\langle \mathbf{G}_i, \mathbf{s}_j \rangle = \mathbf{u}_j[i]$

Problem: Still $O(\lambda^2)$ queries!

IOP for Inner-Products and Proximity



1. \mathbf{P} sends $\mathbf{u} = \text{Enc}_{RS}(\mathbf{s}_1 \parallel \dots \parallel \mathbf{s}_\nu)$ under RS code

2. \mathbf{P} convinces \mathbf{V} that:

1. \mathbf{u} is actually an encoding of $\mathbf{s}_1 \parallel \dots \parallel \mathbf{s}_\nu$:

1. \mathbf{u} is close to D

2. Message in \mathbf{u} is consistent with \mathbf{u}_j 's

2. For each j , $\langle \mathbf{s}_j, \mathbf{g} \rangle = \sigma_j$

Attempt 2:

1. \mathbf{V} samples 4ν random columns \mathbf{u}_j .

2. Per column, she picks $O(1)$ random indices.

3. For each index i , \mathbf{P} proves $\langle \mathbf{G}_i, \mathbf{s}_{j_i} \rangle = \mathbf{u}_{j_i}[i]$

Efficiency: Just $O(4\lambda)$ queries!

Soundness: ??

Soundness

Assume distance of column code is $1/2$.

Claim: Out of 4ν sampled columns \mathbf{u}_j , number of ‘bad’ columns that are far from code is concentrated around average number of bad columns, i.e., 2ν .

Proof: Chernoff Bound asserts that probability that there exist 3ν good columns is negligible.

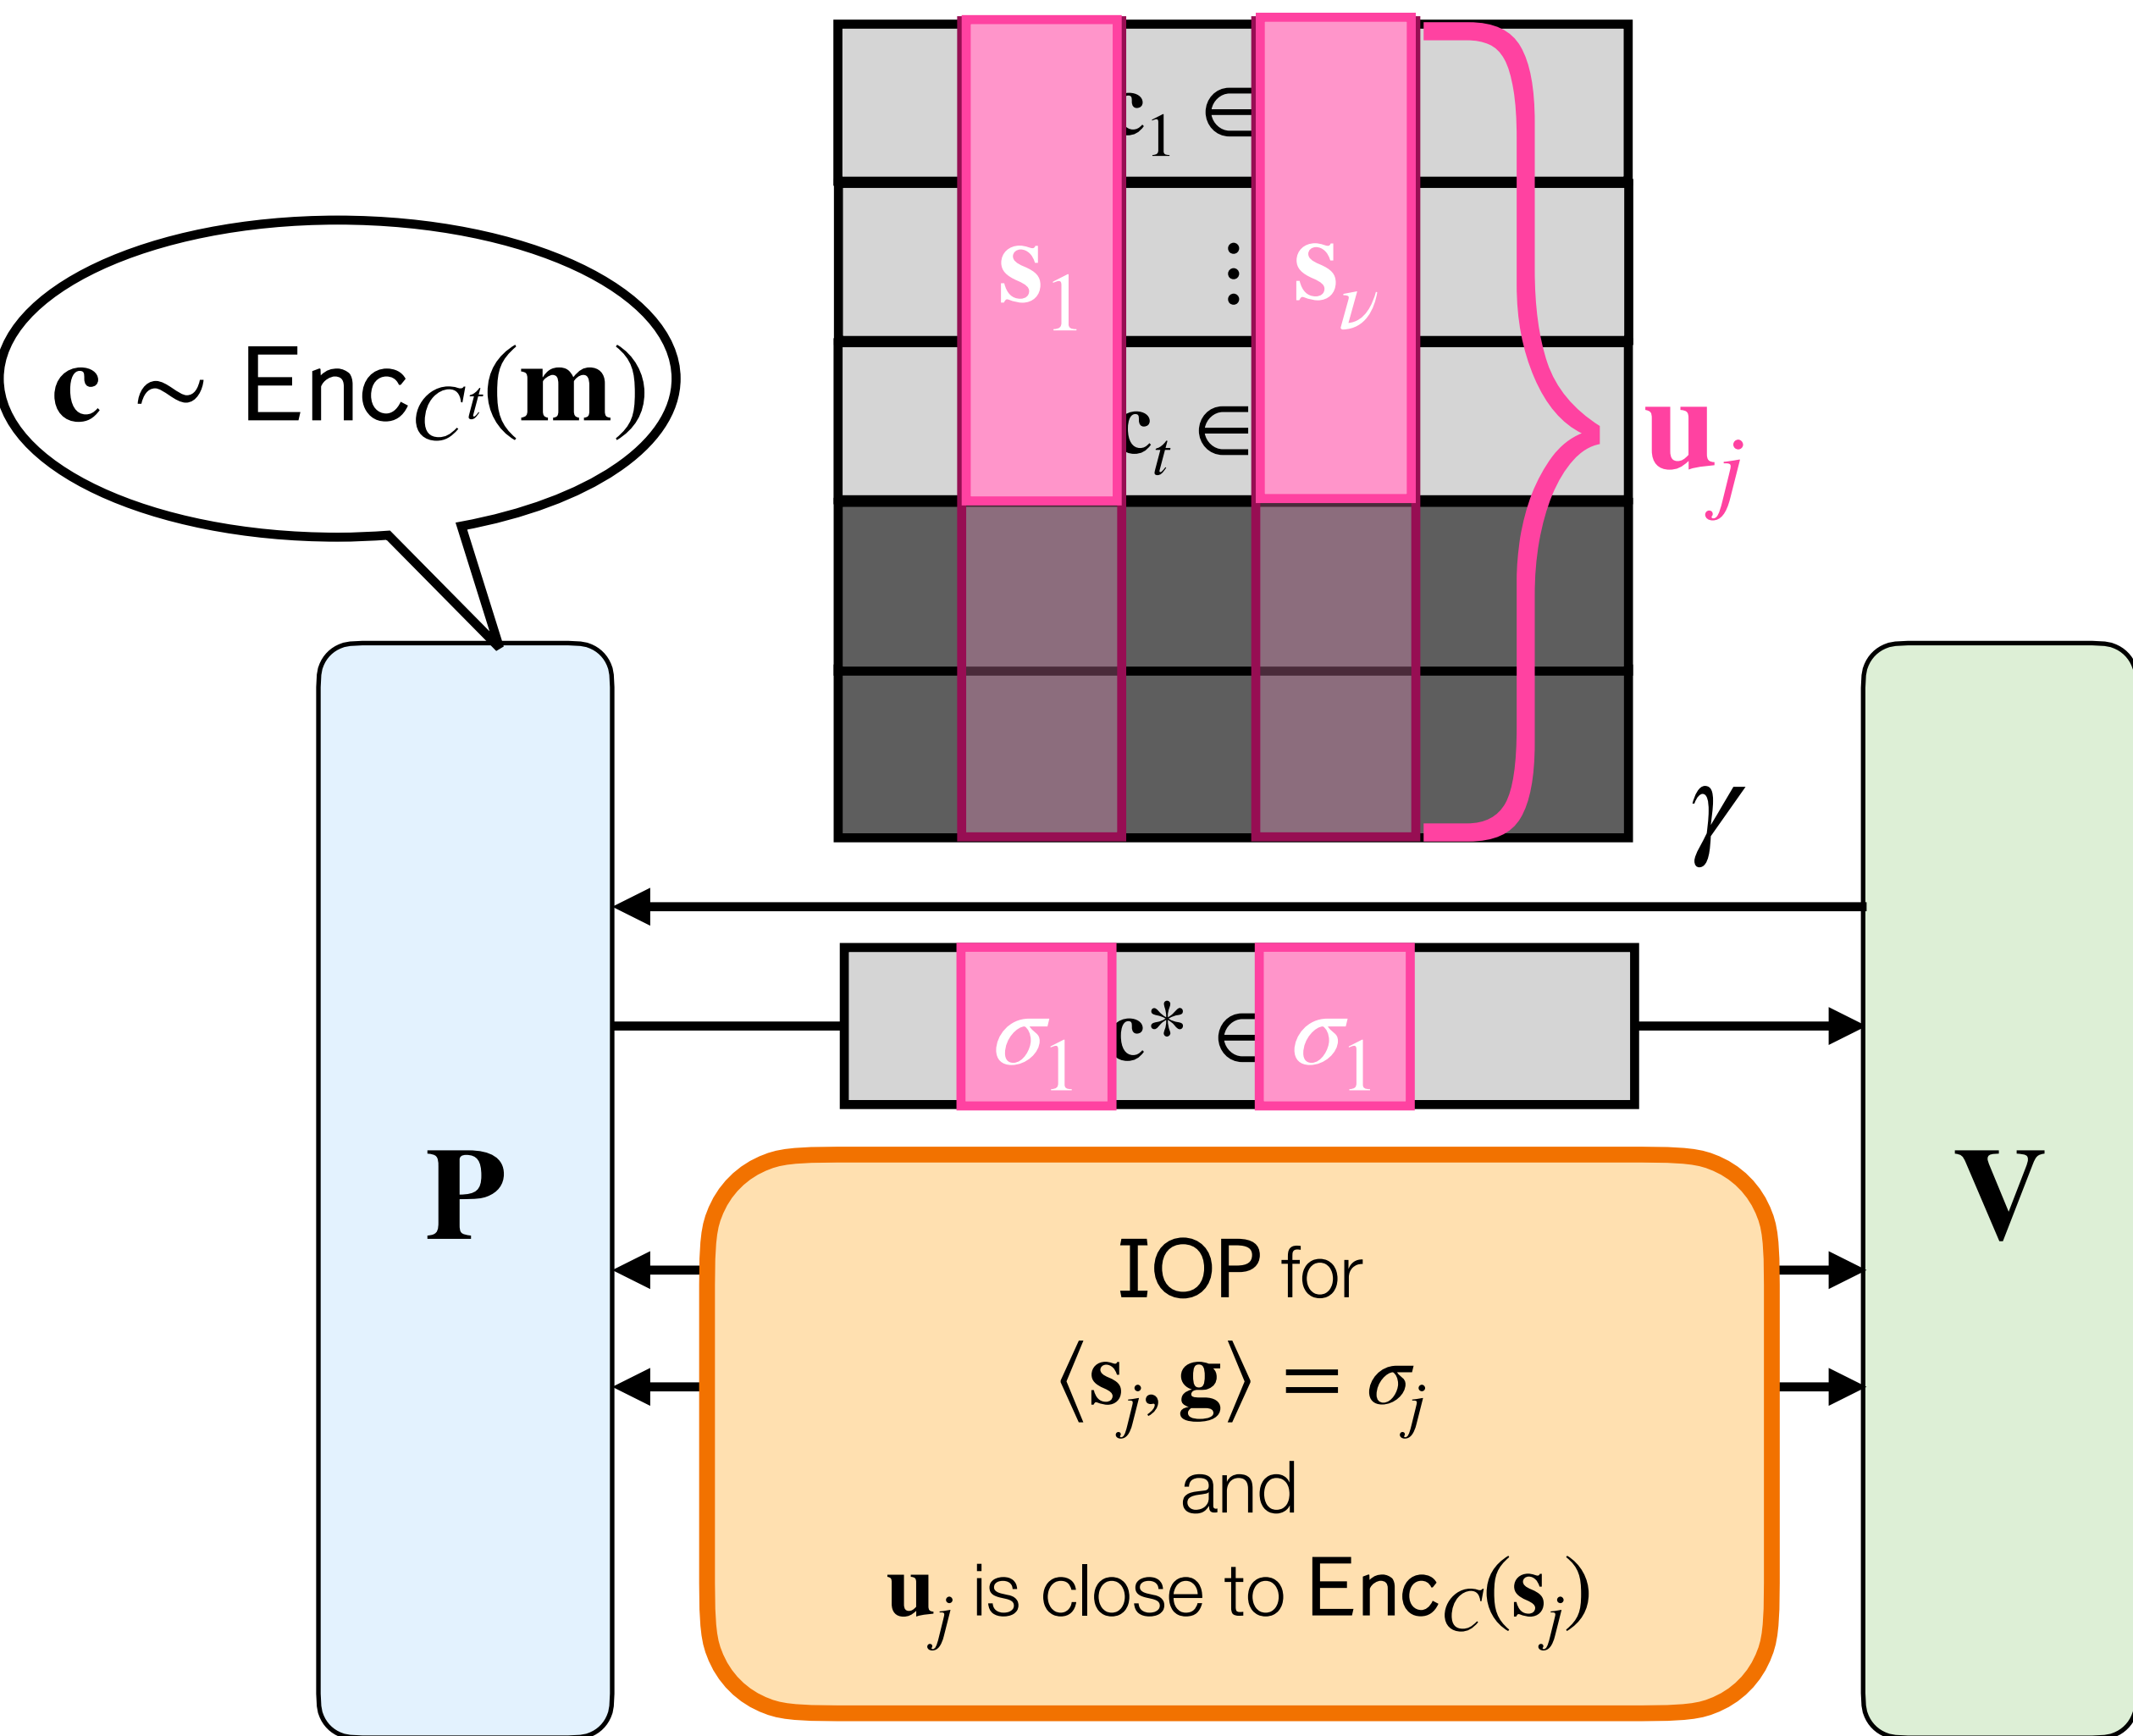
Corollary: With overwhelming probability, at least ν of the \mathbf{u}_j 's are bad.

Now, since \mathbf{V} makes $O(1)$ queries per column, and each bad column is at least $1/2$ -distance away from code, \mathbf{P} gets caught in each bad column with constant probability (say, p).

Combined with corollary, this means that \mathbf{P} avoids detection in each bad column is $(1 - p)^\nu$.

By setting ν appropriately, this is negligible.

Efficiency



Prover time: $T_C + T_{\text{IOP}} = O(T_C)$.

Query complexity: $4\nu + O(\lambda) = O(\lambda)$

