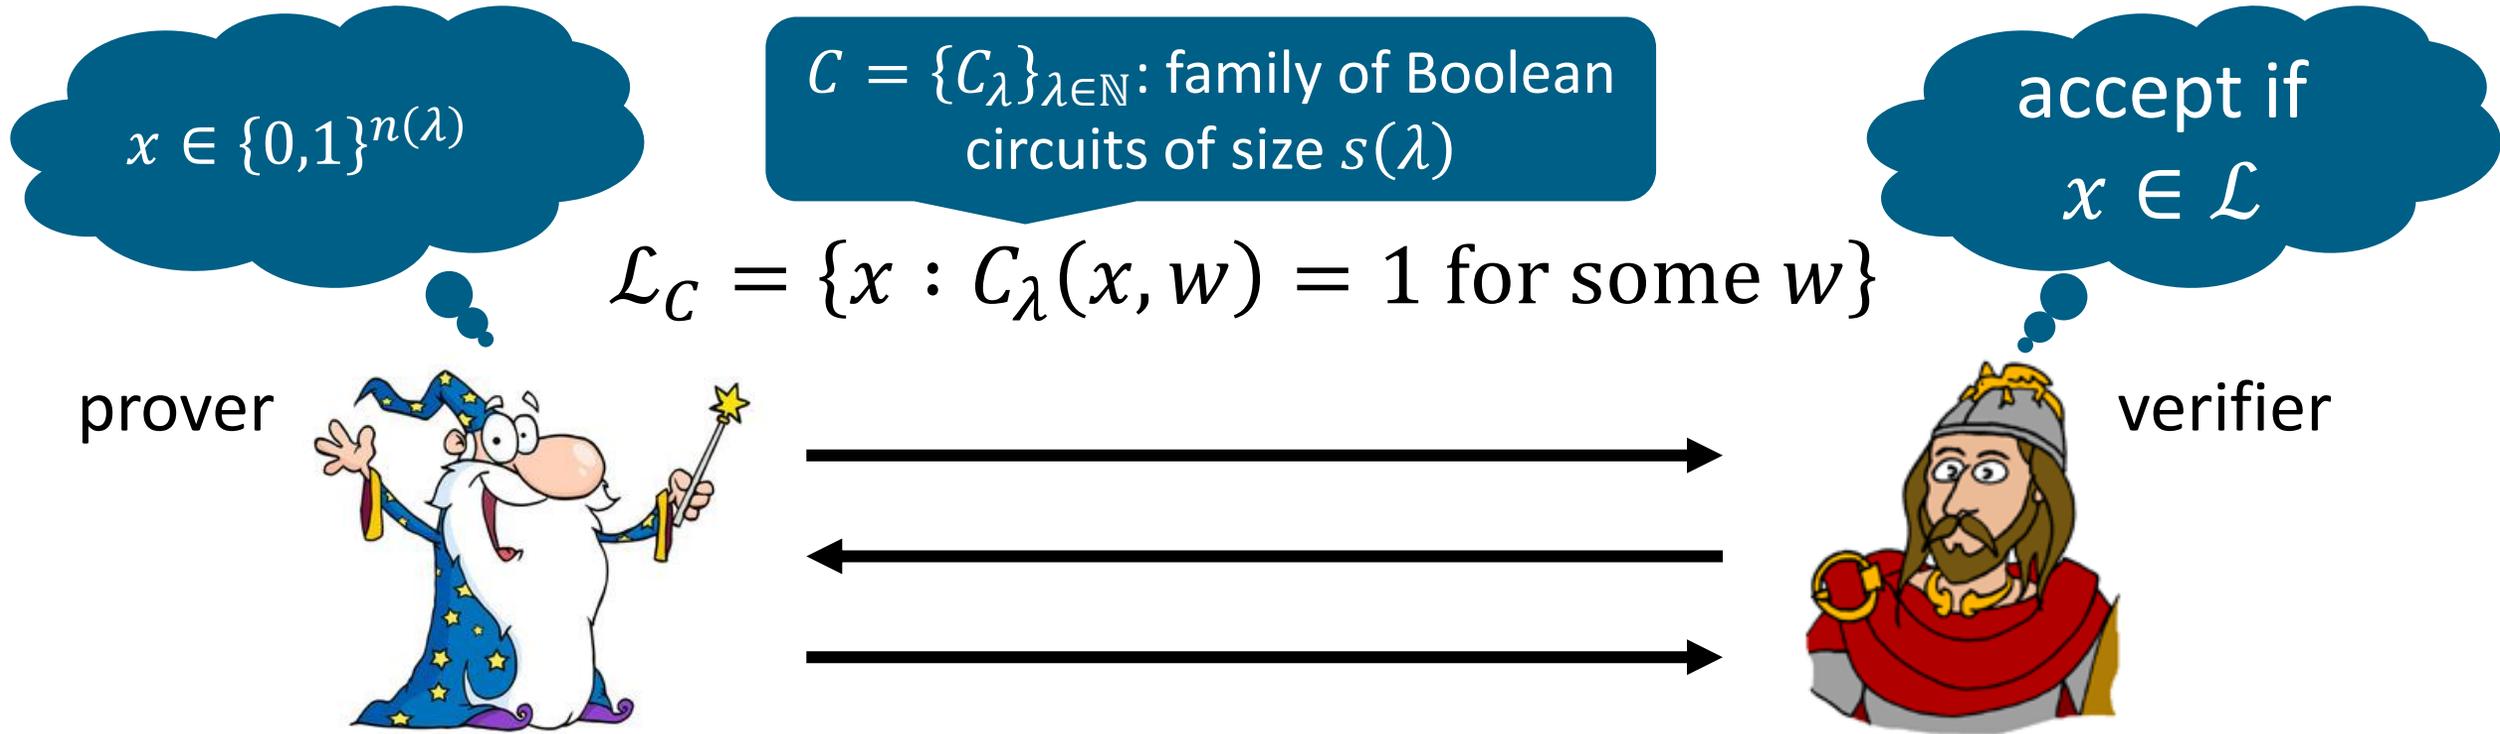


Post-Quantum Designated-Verifier zkSNARKs from Lattices

David Wu
October 2021

Argument Systems

[GMR85]



Completeness:

$$\forall x \in \mathcal{L}_C : \Pr[\langle P(x, w), V(x) \rangle = \text{accept}] = 1$$

“Honest prover convinces honest verifier of true statements”

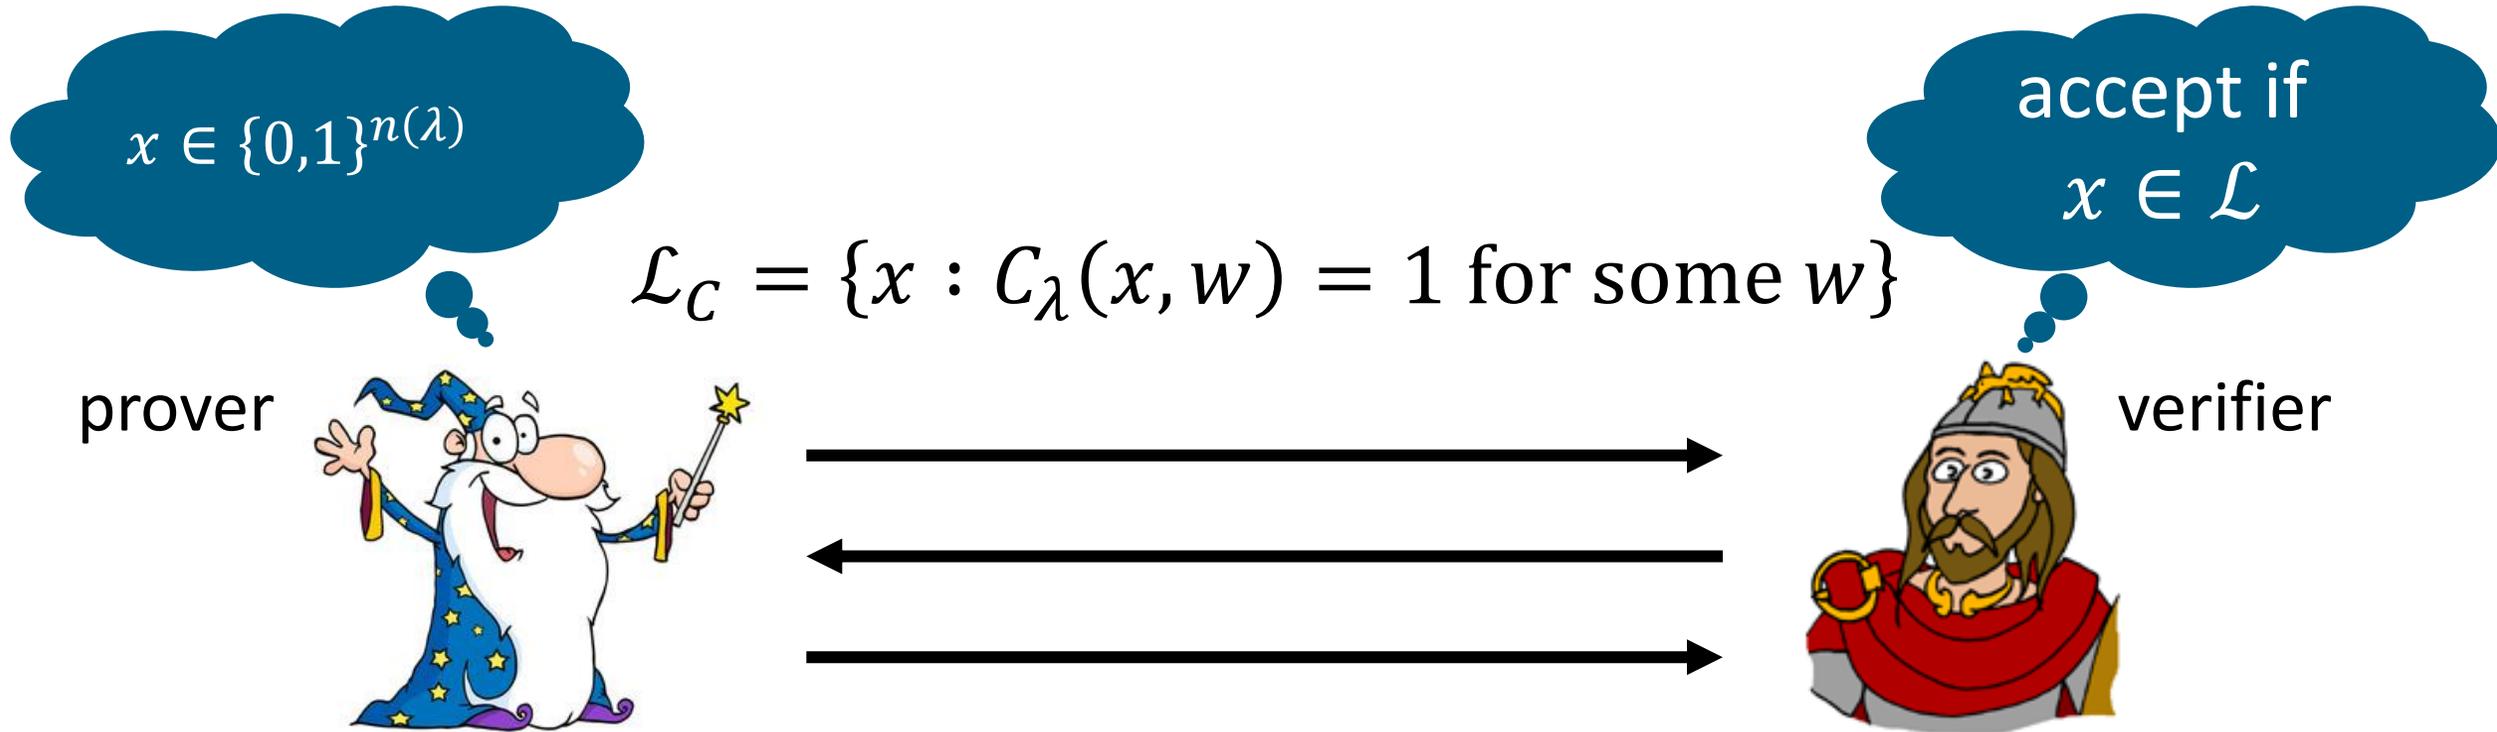
Soundness:

$$\forall x \notin \mathcal{L}_C, \forall \text{ efficient } P^* : \Pr[\langle P^*(1^\lambda, x), V(x) \rangle = \text{accept}] = \text{negl}(\lambda)$$

“Efficient prover cannot convince honest verifier of false”

Argument Systems

[GMR85]



Argument system is **succinct** if:

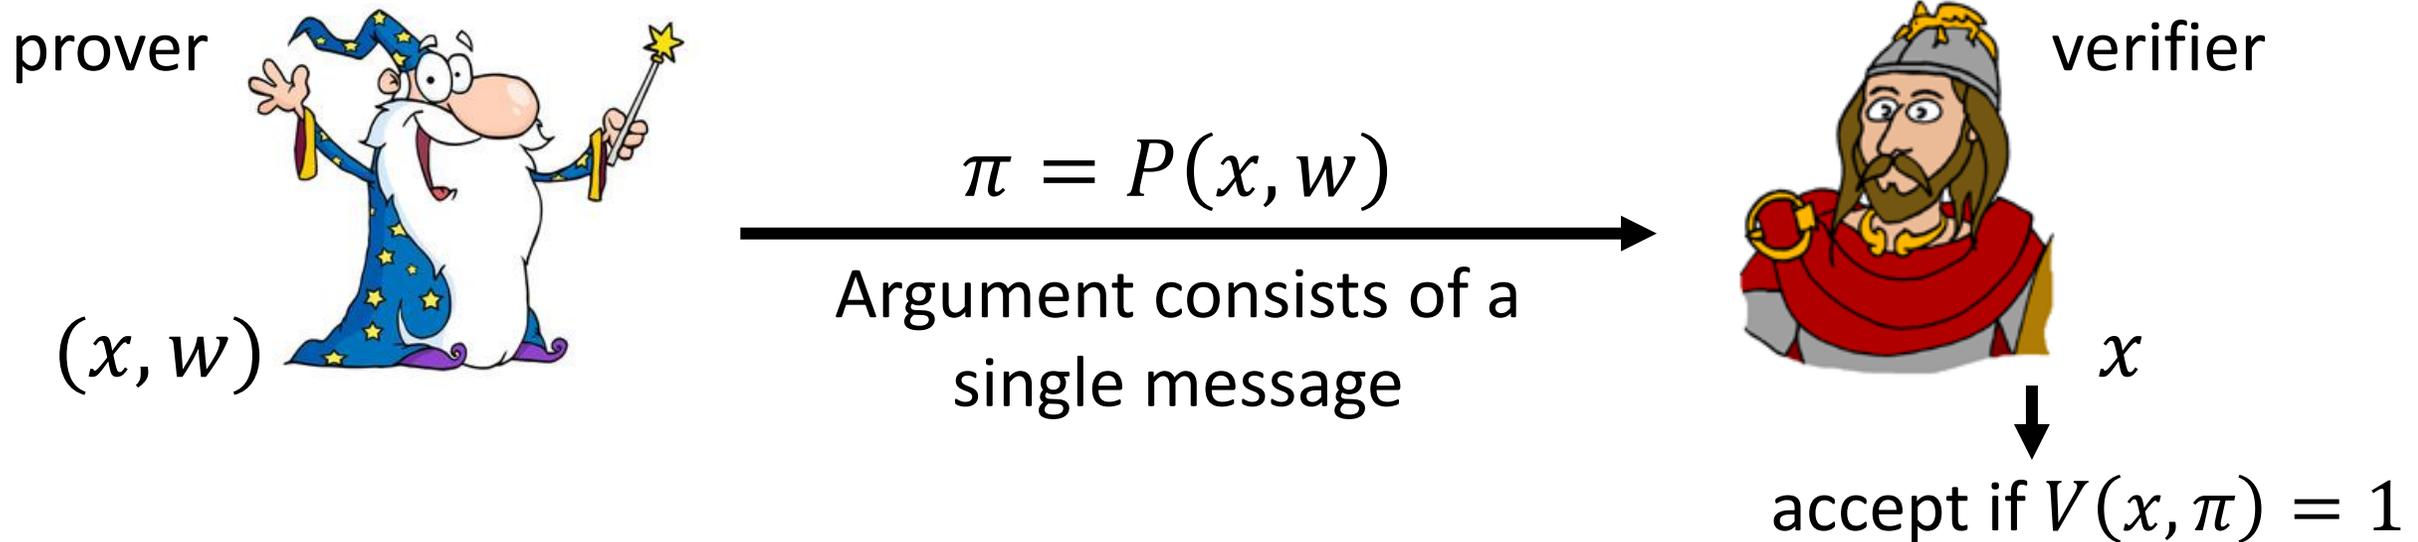
- Prover communication is $\text{poly}(\lambda + \log|C_\lambda|)$
- Running time of V is $\text{poly}(\lambda + |x| + \log|C_\lambda|)$

Both must be smaller
than classic NP
verification

Succinct Non-Interactive Arguments (SNARGs)

[Kil92, Mic00, GW11]

$$\mathcal{L}_C = \{x : C_\lambda(x, w) = 1 \text{ for some } w\}$$



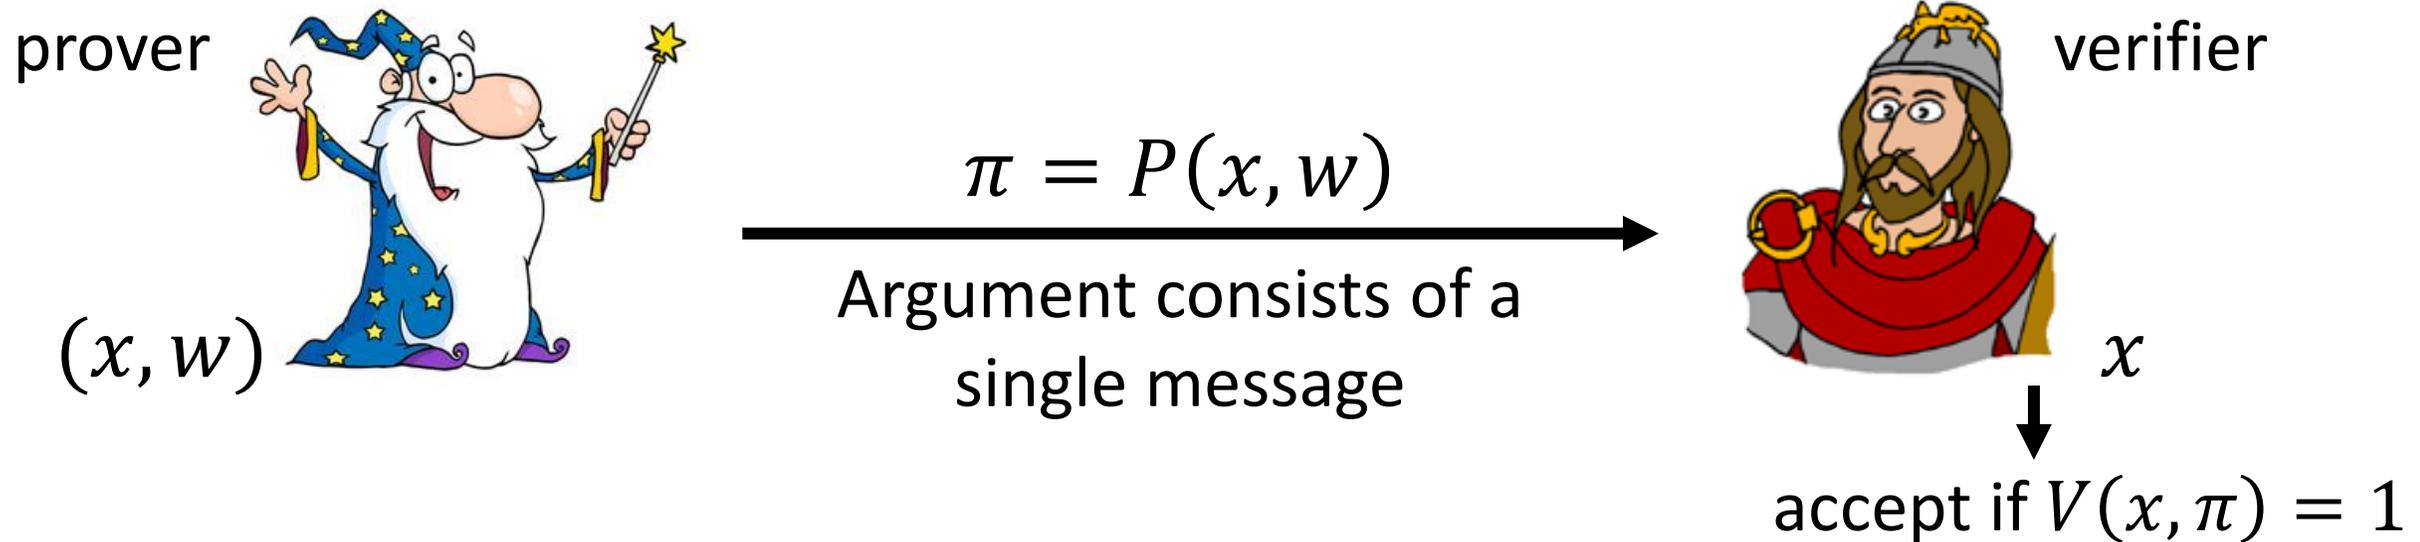
Additional properties of interest:

- **Proof of knowledge**: succinct non-interactive argument of knowledge (SNARK):
“There exists an efficient extractor that can recover a witness from any prover that convinces an honest verifier”

Succinct Non-Interactive Arguments (SNARGs)

[Kil92, Mic00, GW11]

$$\mathcal{L}_C = \{x : C_\lambda(x, w) = 1 \text{ for some } w\}$$



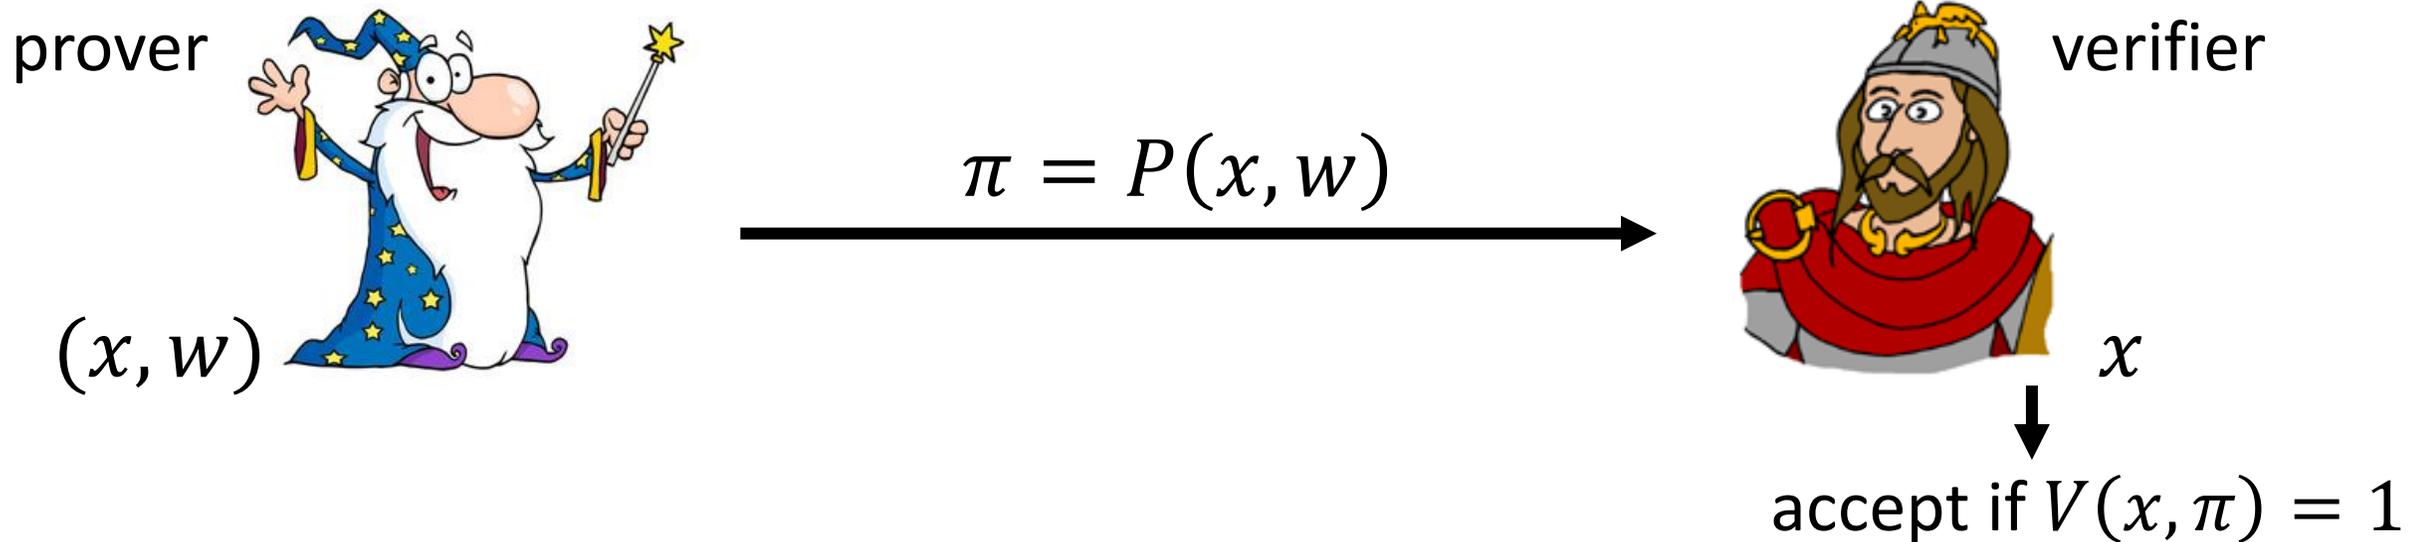
Additional properties of interest:

- **Zero-knowledge:** “Proof does not leak information about the prover’s witness”
- **zkSNARK:** zero-knowledge succinct non-interactive argument of knowledge

Succinct Non-Interactive Arguments (SNARGs)

[Kil92, Mic00, GW11]

$$\mathcal{L}_C = \{x : C_\lambda(x, w) = 1 \text{ for some } w\}$$

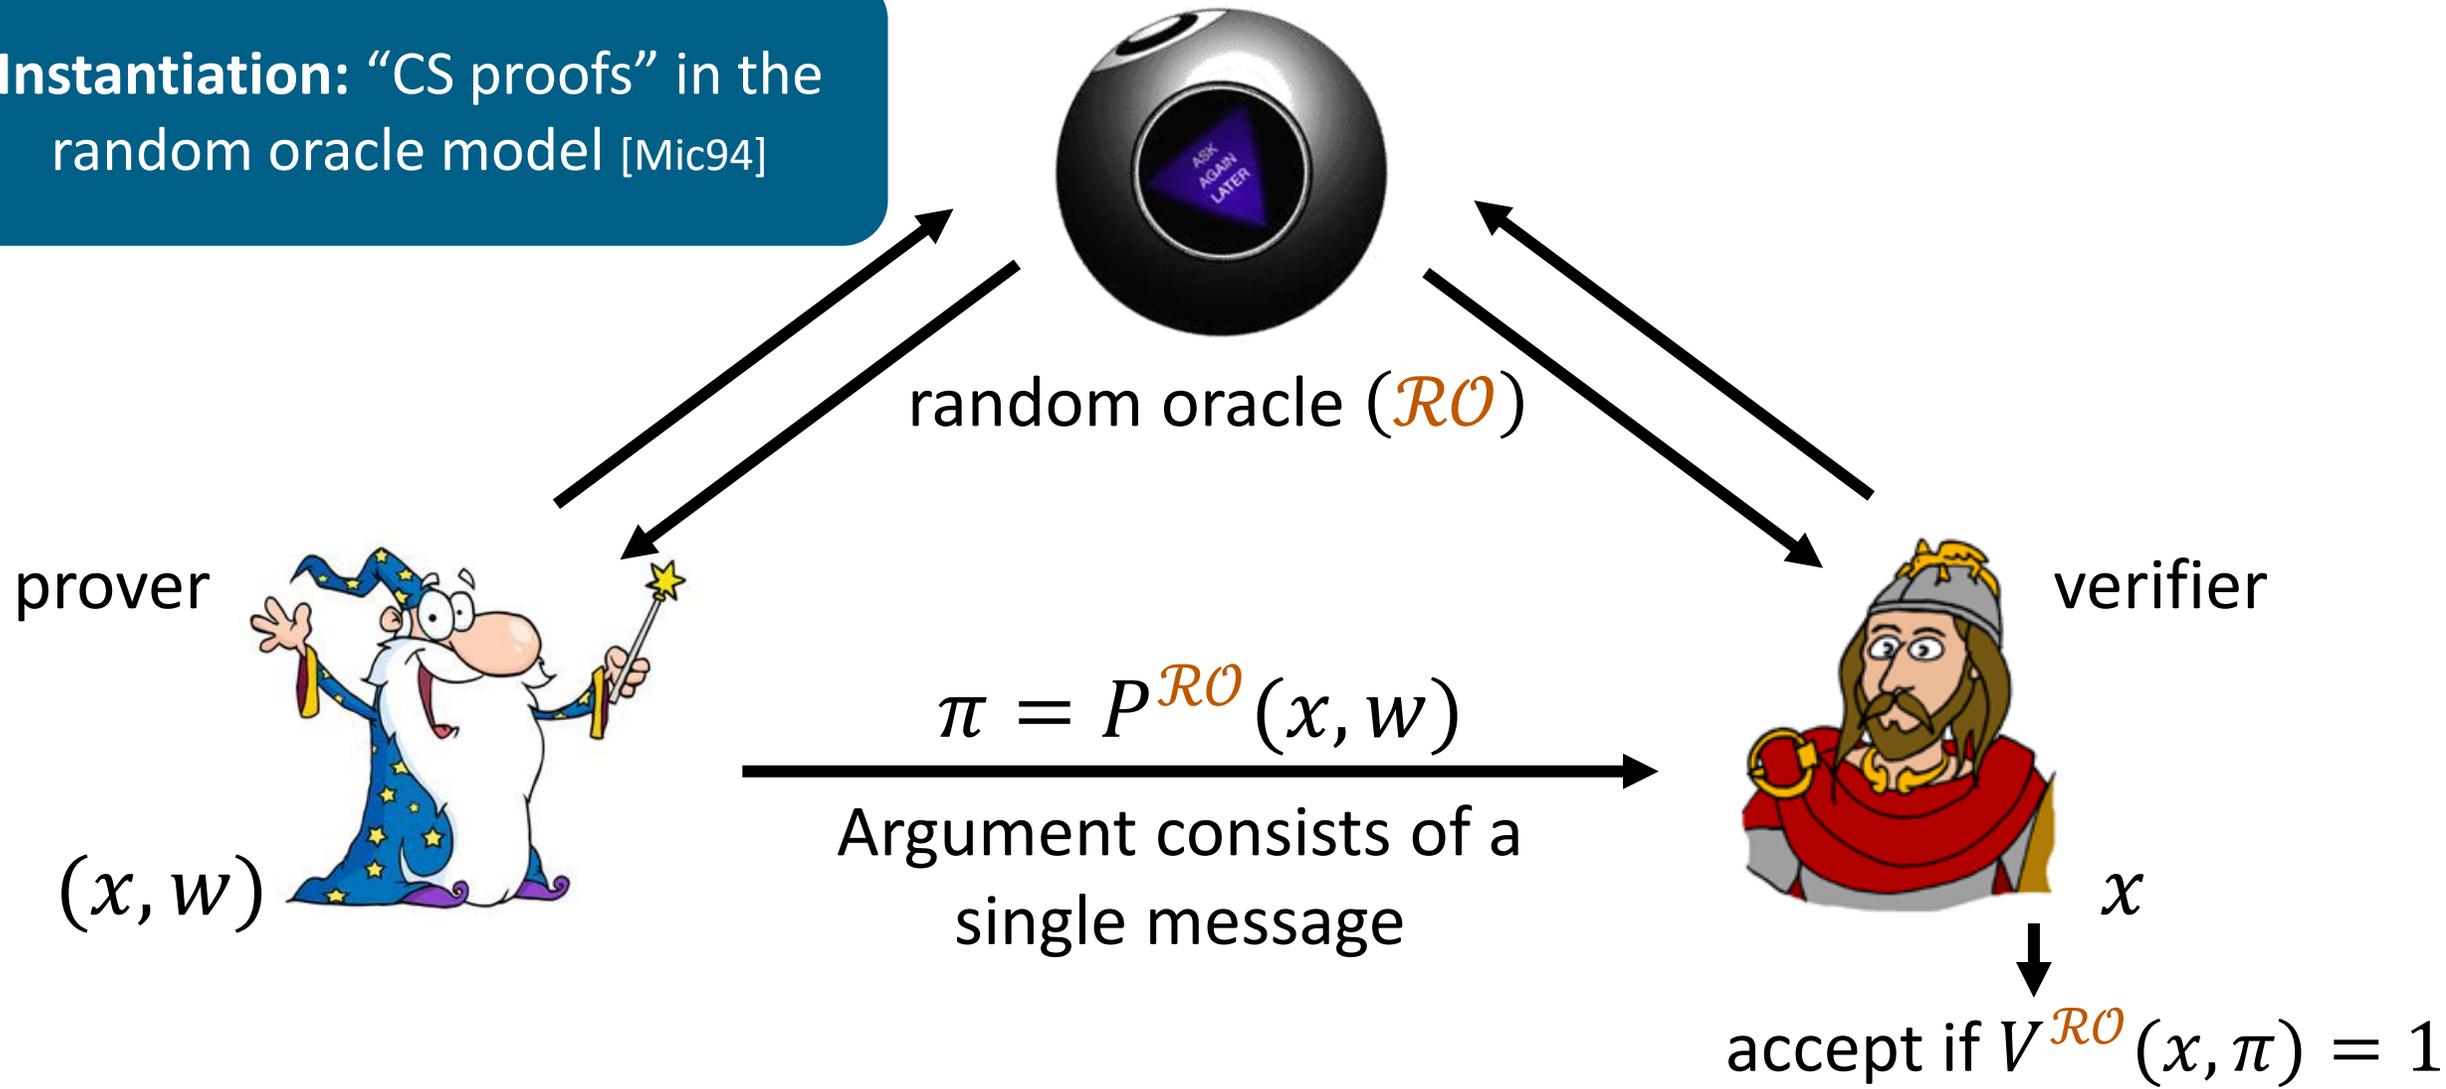


For general NP languages, SNARGs are unlikely to exist in standard model [BP04, Wee05]

Succinct Non-Interactive Arguments (SNARGs)

[Kil92, Mic00, GW11]

Instantiation: "CS proofs" in the random oracle model [Mic94]



Succinct Non-Interactive Arguments (SNARGs)

[Kil92, Mic00, GW11]

Preprocessing SNARGs:
allow “expensive” setup

Setup(1^λ)



common reference
string (CRS)

verification
state



Can consider publicly-
verifiable and secretly-
verifiable SNARGs

prover



(x, w)

$$\pi = P(\sigma, x, w)$$



verifier



x

accept if $V(\tau, x, \pi) = 1$

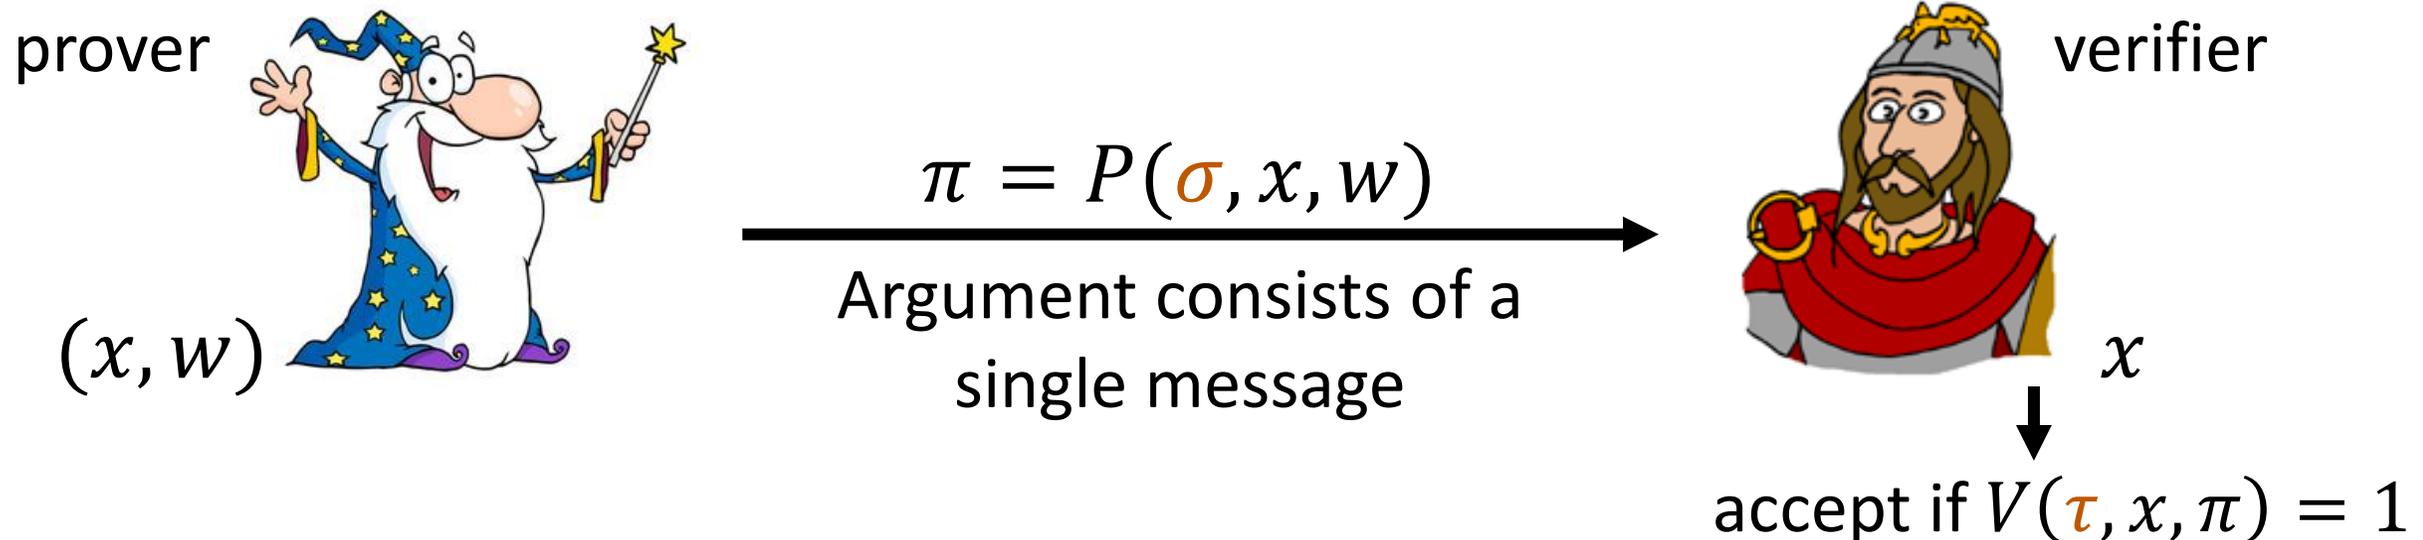
Succinct Non-Interactive Arguments (SNARGs)

[Kil92, Mic00, GW11]

Very active area of research (encompassing both theory and practice):

PHGR13, BCI⁺13, BCC⁺16, Gro16, ZGK⁺17, AHIV17, WTS⁺18, GMNO18, BBB⁺18, BBHR19, BCR⁺19, XZZ⁺19, LM19, CHM⁺20, BFS20, SL20, Set20, COS20, CY21, GNS21, GMN21, GLS⁺21, and *many, many more...*

This talk: post-quantum constructions (specifically, from lattice-based assumptions)



zkSNARK Constructions (with Implementation)

Construction	Prover Complexity	Proof Size		Assumption
		Asymptotic	Concrete	
[Gro16]	$N \log N$	1	128 bytes	Pairings
Marlin [CHM ⁺ 20]	$N \log N$	1	704 bytes	Pairings
Xiphos [SL20]	N	$\log N$	61 KB	Pairings <i>Pre-Quantum</i>
Fractal [COS20]	$N \log N$	$\log^2 N$	215 KB	Random Oracle
STARK [BBHR19]	$N \text{ polylog } N$	$\log^2 N$	127 KB*	Random Oracle
[GMNO18] [†]	$N \log N$	1	640 KB	Lattices <i>Post-Quantum</i>

[†]designated-verifier

Focus is on constructions with a *succinct* verifier

*for a structured computation

N : size of NP relation being verified ($N \approx 2^{20}$ for concrete values)

Asymptotic metrics are given up to $\text{poly}(\lambda)$ factors (for a security parameter λ)

zkSNARK Constructions (with Implementation)

Construction	Prover Complexity	Proof Size		Assumption
		Asymptotic	Concrete	
[Gro16]	$N \log N$	1	128 bytes	Pairings
Marlin [CHM ⁺ 20]	$N \log N$	1	704 bytes	Pairings
Xiphos [SL20]	N	$\log N$	61 KB	Pairings <i>Pre-Quantum</i>
Fractal [COS20]	$N \log N$	$\log^2 N$	215 KB	Random Oracle
STARK [BBHR19]	$N \text{ polylog } N$	$\log^2 N$	127 KB*	Random Oracle
[GMNO18] [†]	$N \log N$	1	640 KB	Lattices <i>Post-Quantum</i>

1000× gap between size of pre-quantum zkSNARKs and post-quantum ones

This talk: constructing shorter post-quantum zkSNARKs (via lattice-based assumptions)

zkSNARK Constructions (with Implementation)

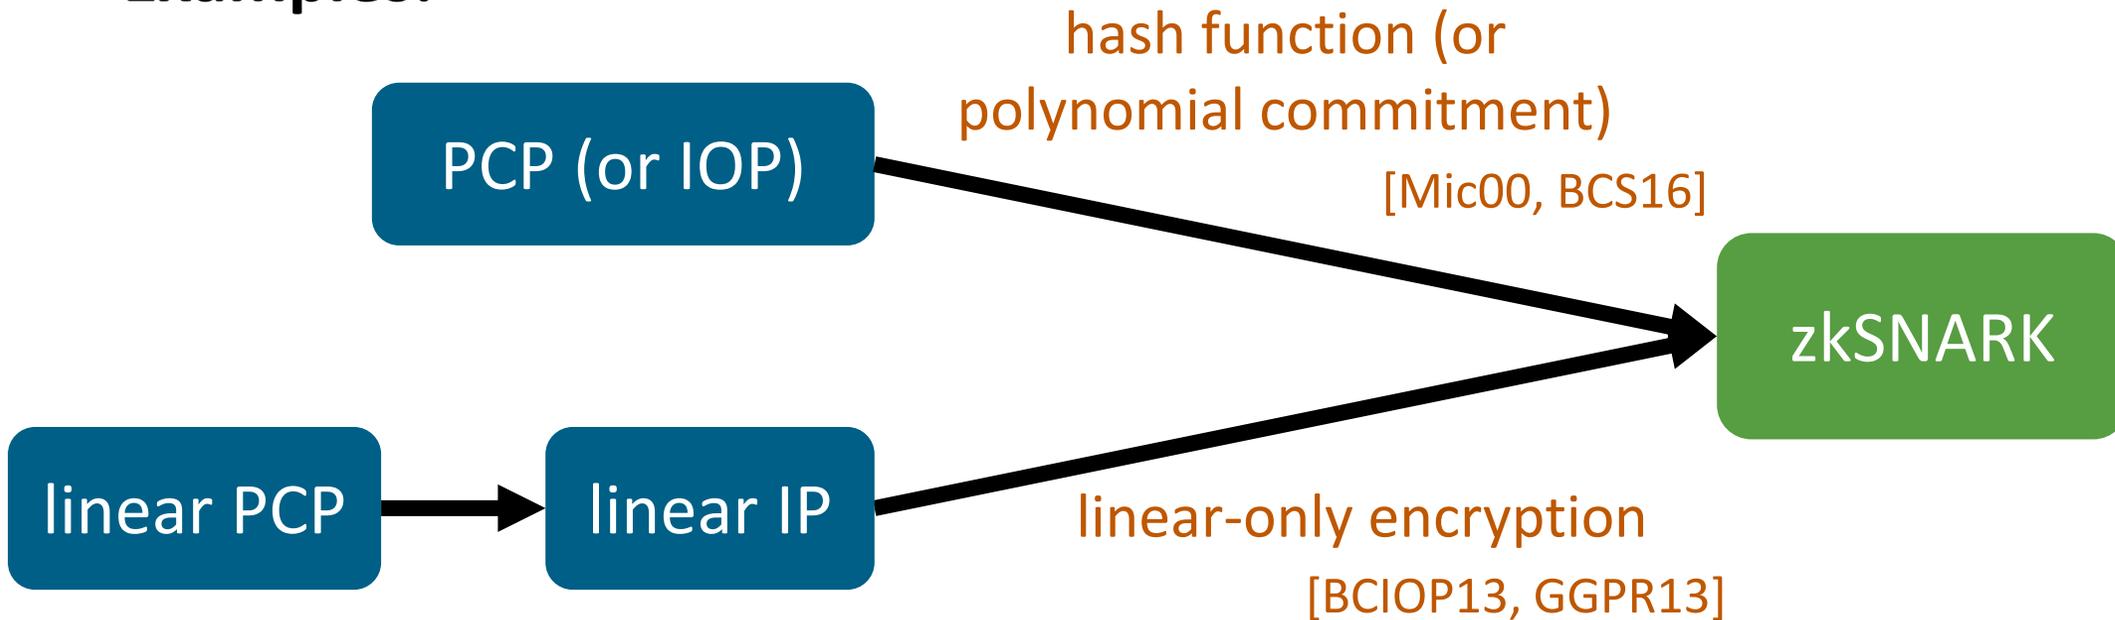
Construction	Prover Complexity	Proof Size		Assumption
		Asymptotic	Concrete	
[Gro16]	$N \log N$	1	128 bytes	Pairings
Marlin [CHM ⁺ 20]	$N \log N$	1	704 bytes	Pairings
Xiphos [SL20]	N	$\log N$	61 KB	Pairings <i>Pre-Quantum</i>
Fractal [COS20]	$N \log N$	$\log^2 N$	215 KB	Random Oracle
STARK [BBHR19]	$N \text{ polylog } N$	$\log^2 N$	127 KB*	Random Oracle
[GMNO18] [†]	$N \log N$	1	640 KB	Lattices
This work	$N \log N$	1	16 KB	Lattices <i>Post-Quantum</i>

- $\approx 10\times$ shorter proofs compared to previous post-quantum zkSNARKs for general NP relations
- Prover and verifier are concretely faster compared to most succinct pre-quantum construction [Gro16]
- Construction is designated-verifier (need secret key to check proofs) and has long CRS

Construction Overview

Follows the classic approach of combining an [information-theoretic](#) proof system (for NP) with a [cryptographic](#) compiler

Examples:

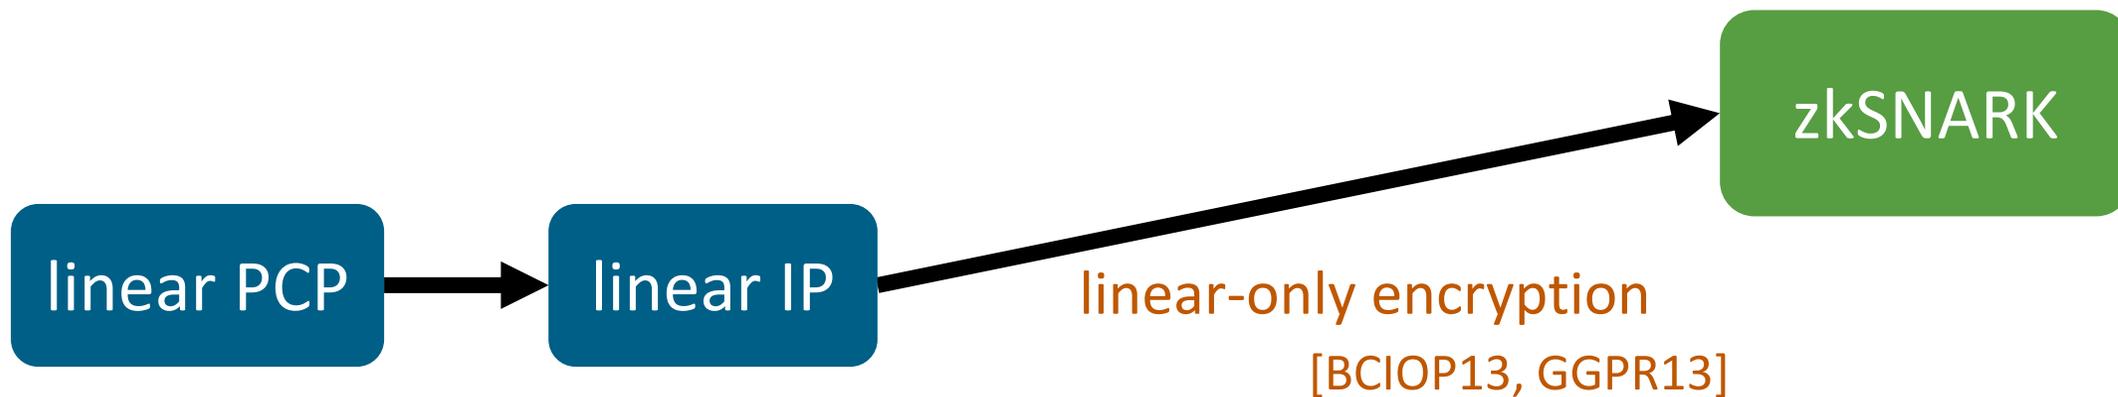


Construction Overview

Follows the classic approach of combining an [information-theoretic](#) proof system (for NP) with a [cryptographic](#) compiler

Starting point: the [BCIOP13] compiler from linear PCPs to zkSNARKs

- Yields the most succinct pre-quantum zkSNARKs [GGPR13, Gro16]
- Basis of several lattice-based zkSNARKs [BISW17, GMNO18]



Linear Probabilistically-Checkable Proofs (LPCPs)

[IKO07]

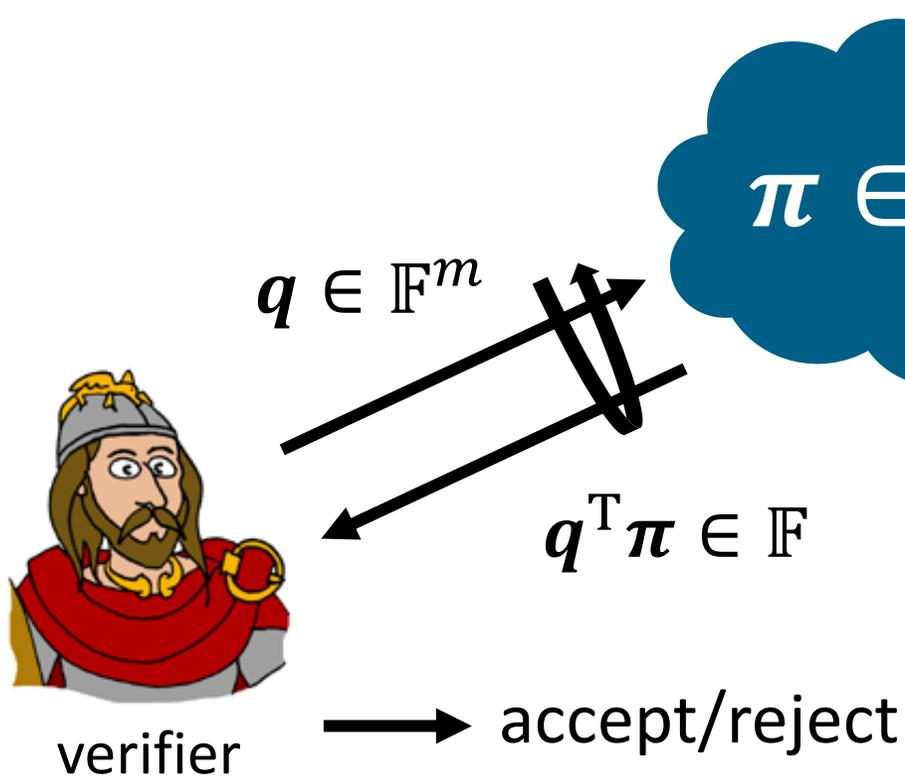
(x, w)



$\pi \in \mathbb{F}^m$

linear PCP

“encoding” of statement/witness



- Verifier given oracle access to a *linear* function $\pi \in \mathbb{F}^m$
- Several instantiations:
 - 3-query LPCP based on the Walsh-Hadamard code: $m = O(|C|^2)$ [ALMSS92]
 - 4-query LPCP based on quadratic arithmetic programs: $m = O(|C|)$ [GGPR13]

Linear Probabilistically-Checkable Proofs (LPCPs)

[IKO07]

(x, w)



$\pi \in \mathbb{F}^m$

linear PCP

“encoding” of statement/witness

$\pi \in \mathbb{F}^m$

$q \in \mathbb{F}^m$

$q^T \pi \in \mathbb{F}$



verifier

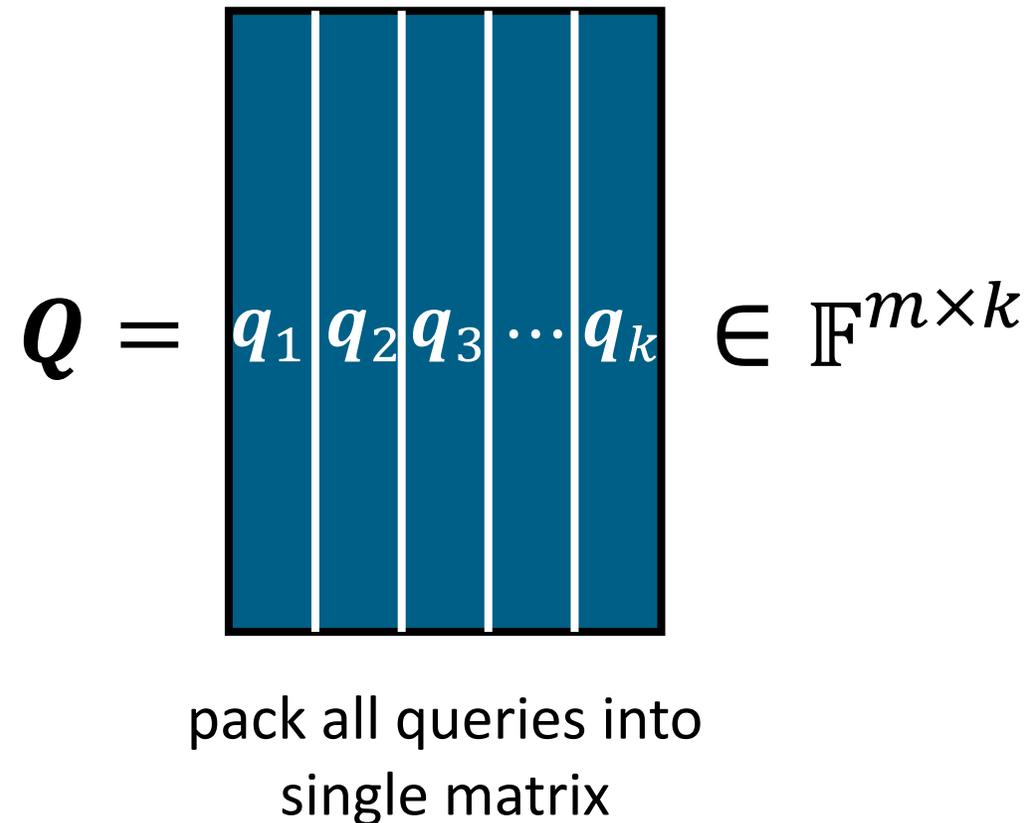
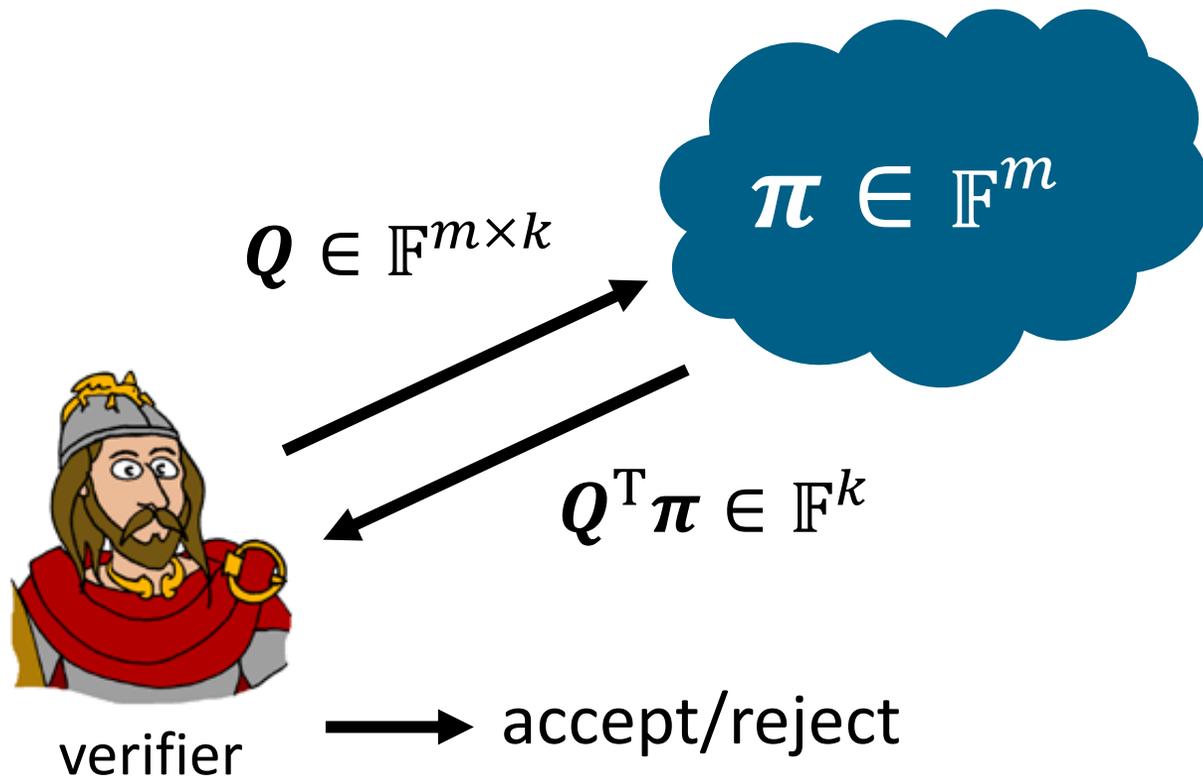
→ accept/reject

Oftentimes, verifier is *oblivious*:
the queries q do not depend on
the statement x

Linear Probabilistically-Checkable Proofs (LPCPs)

[IKO07]

Equivalent view (if verifier is oblivious):



From Linear PCPs to Preprocessing SNARGs

[BCIOP13]

Oblivious verifier can “commit”
to its queries ahead of time



$$Q = \underbrace{\begin{matrix} | & | & | & \cdots & | \\ q_1 & q_2 & q_3 & \cdots & q_k \end{matrix}}_{\text{part of the CRS}}$$

part of the CRS



Honest prover takes (x, w) and constructs
linear PCP $\pi \in \mathbb{F}^m$ and
computes $Q^T \pi$

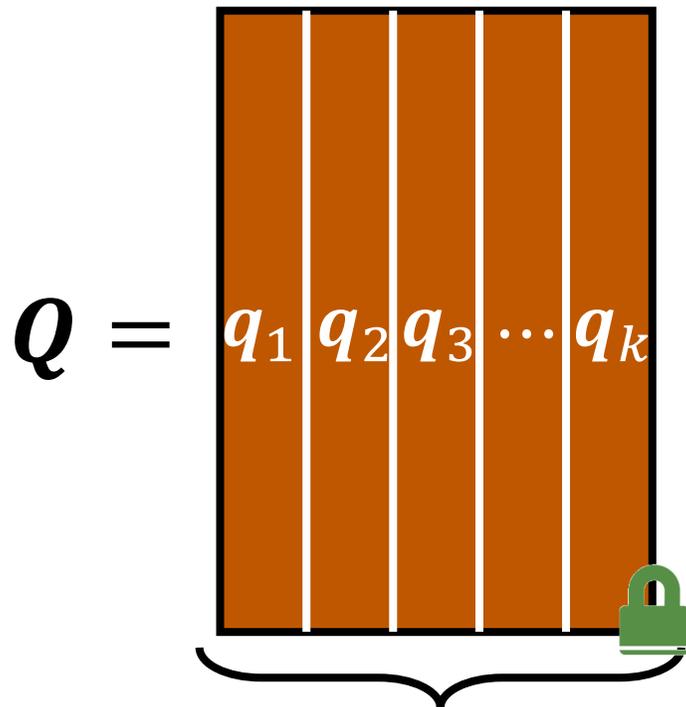
Two problems:

- Malicious prover can choose π based on queries
- Malicious prover can apply different π to the different columns of Q

From Linear PCPs to Preprocessing SNARGs

[BCIOP13]

Oblivious verifier can “commit”
to its queries ahead of time



part of the CRS



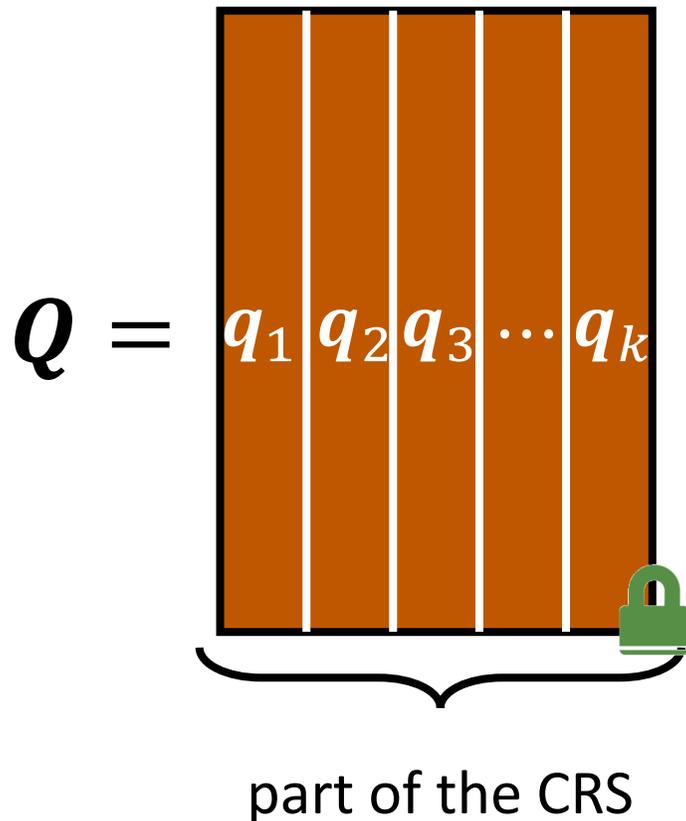
Honest prover takes (x, w) and constructs
linear PCP $\pi \in \mathbb{F}^m$ and
computes $Q^T \pi$

Step 1: Encrypt elements of Q using
additively homomorphic encryption scheme

From Linear PCPs to Preprocessing SNARGs

[BCIOP13]

Oblivious verifier can “commit”
to its queries ahead of time



Honest prover takes (x, w) and constructs
linear PCP $\pi \in \mathbb{F}^m$ and
computes $Q^T \pi$

homomorphic
evaluation



From Linear PCPs to Preprocessing SNARGs

[BCIOP13]



Designated-verifier SNARK:
decryption key needed to verify

If LPCP verification can be performed directly on ciphertexts (e.g., with pairing-based instantiations), then SNARK is **publicly-verifiable**



Honest prover takes (x, w) and constructs linear PCP $\pi \in \mathbb{F}^m$ and computes $Q^T \pi$

homomorphic evaluation

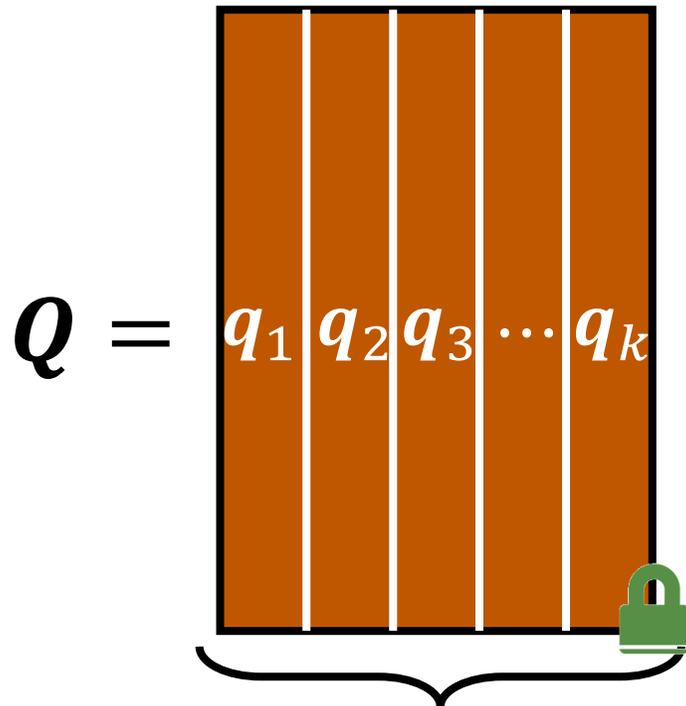


Verifier decrypts to learn $Q^T \pi$ and runs linear PCP decision procedure

From Linear PCPs to Preprocessing SNARGs

[BCIOP13]

Oblivious verifier can “commit”
to its queries ahead of time



part of the CRS



Honest prover takes (x, w) and constructs
linear PCP $\pi \in \mathbb{F}^m$ and
computes $Q^T \pi$

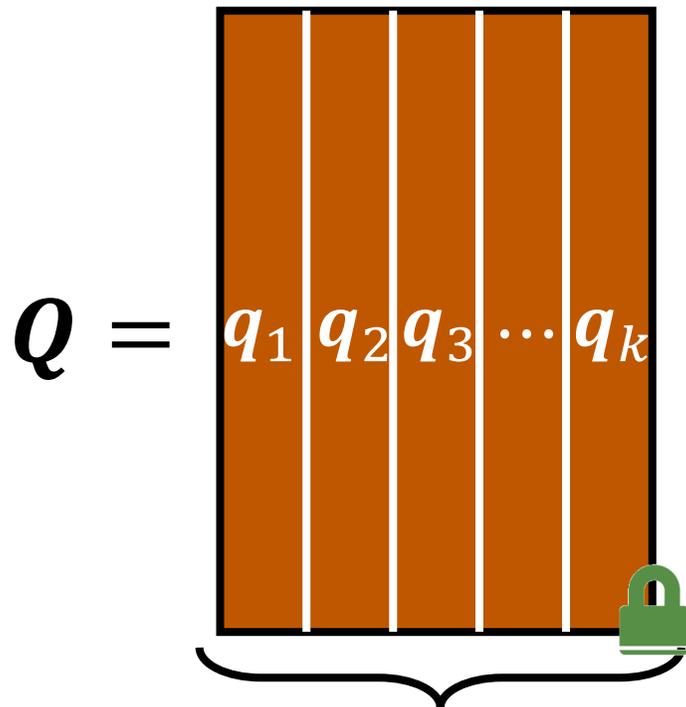
Two problems:

- Malicious prover can choose π based on queries
- Malicious prover can apply different π to the different columns of Q

From Linear PCPs to Preprocessing SNARGs

[BCIOP13]

Oblivious verifier can “commit”
to its queries ahead of time



part of the CRS



Honest prover takes (x, w) and constructs
linear PCP $\pi \in \mathbb{F}^m$ and
computes $Q^T \pi$

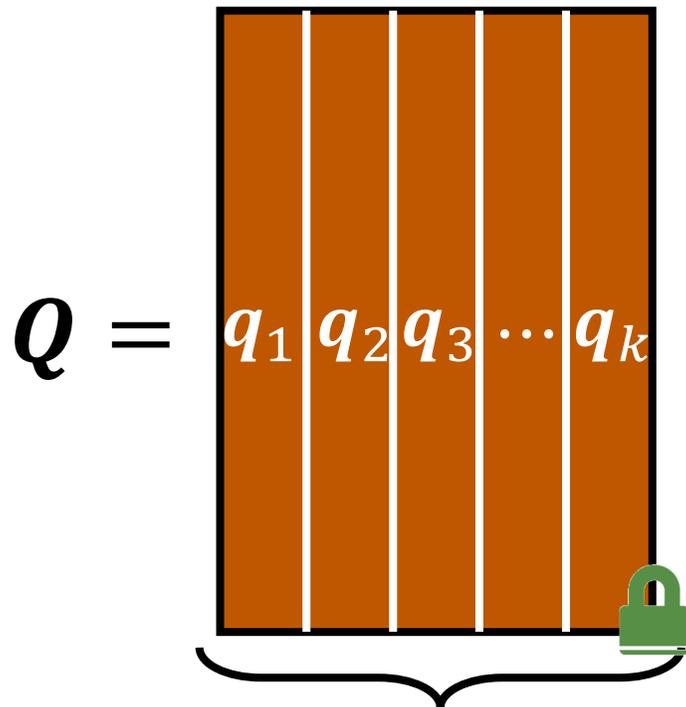
[BCIOP13] approach:

- Add a linear consistency check and view construction as a linear IP (LIP)
- Encrypt the LIP queries using a “linear-only” encryption scheme

From Linear PCPs to Preprocessing SNARGs

[BCIOP13]

Oblivious verifier can “commit”
to its queries ahead of time



part of the CRS



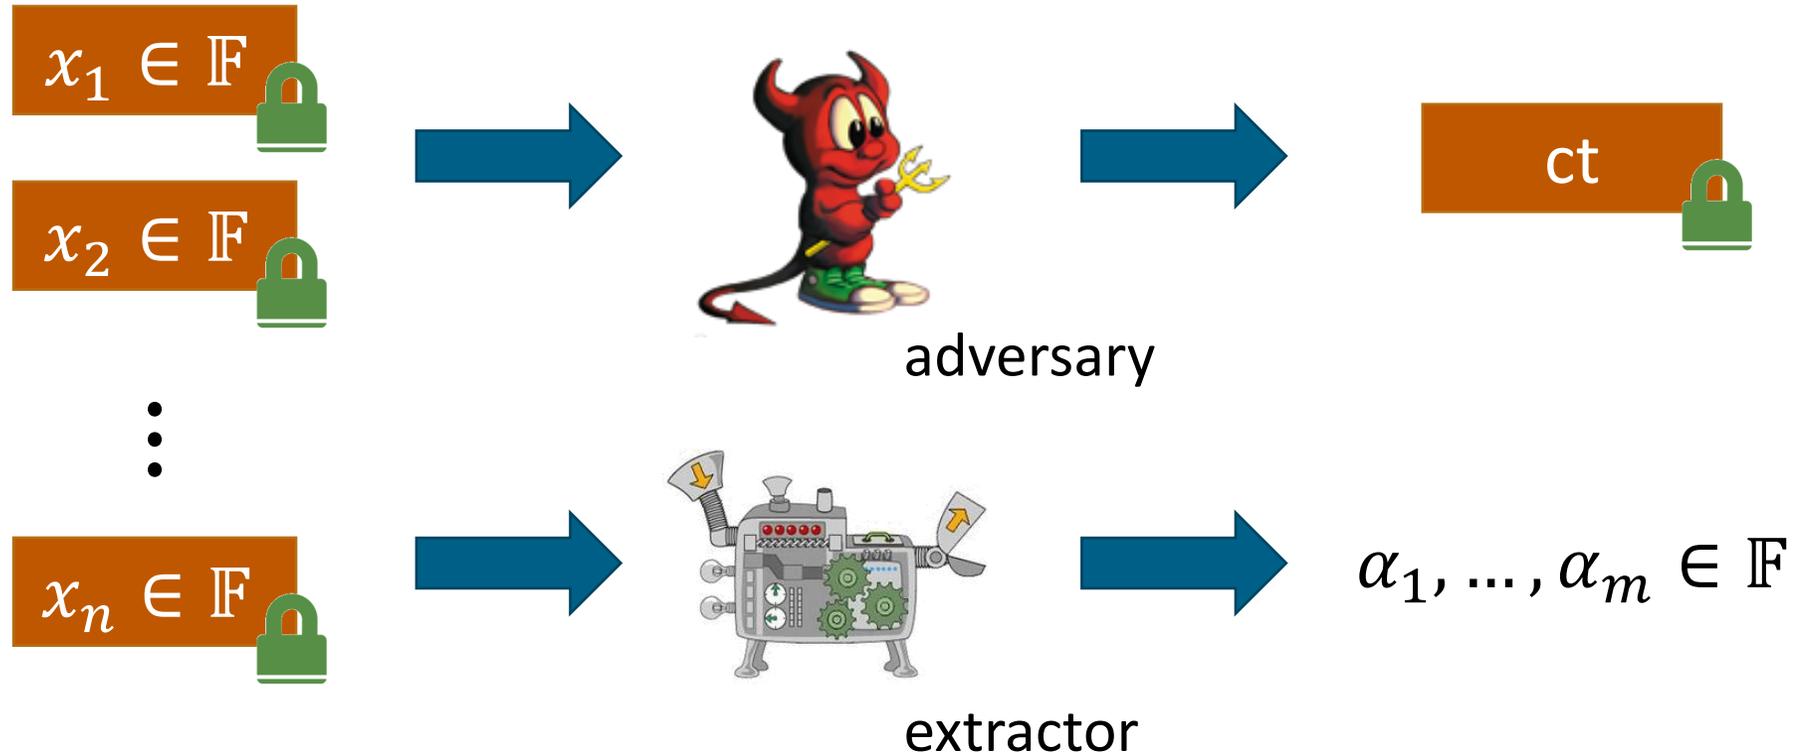
Honest prover takes
 (x, w) and constructs
linear PCP $\pi \in \mathbb{F}^m$ and
computes $Q^T \pi$

Intuitively: an encryption scheme that
only supports additive homomorphism

- Encrypt the LIP queries using a “linear-only” encryption scheme

Linear-Only Encryption

[BCIOP13]



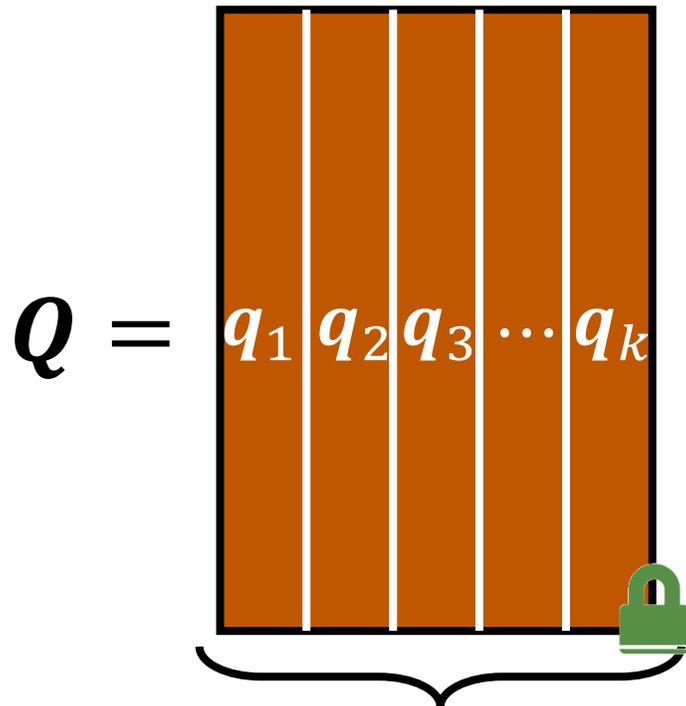
Requirement: If $\text{Decrypt}(\text{sk}, \text{ct}) \neq \perp$, then $\text{Decrypt}(\text{sk}, \text{ct}) = \sum_{i \in [n]} \alpha_i x_i$

Intuition: adversary's strategy can be "explained" by a linear function

From Linear PCPs to Preprocessing SNARGs

[BCIOP13]

Oblivious verifier can “commit”
to its queries ahead of time



part of the CRS



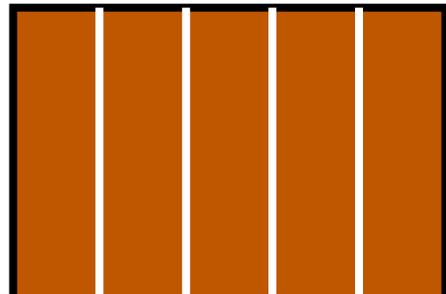
Honest prover takes (x, w) and constructs
linear PCP $\pi \in \mathbb{F}^m$ and
computes $Q^T \pi$

All adversarial strategies can be explained by
a linear function of the encrypted query
components \Rightarrow soundness can now be based
on the soundness of the linear PCP

From Linear PCPs to Preprocessing SNARGs

[BCIOP13]

Oblivious verifier can “commit”
to its queries ahead of time



Honest prover takes
 (x, w) and constructs
linear PCP $\pi \in \mathbb{F}^m$ and
computes $Q^T \pi$

For zero-knowledge, require that LPCP is
(honest-verifier) ZK and encryption scheme
is circuit private (hides linear combination)

Rest of this talk: will not focus on ZK

part of the CRS

All adversarial strategies can be explained by
a linear function of the encrypted query
components \Rightarrow soundness can now be based
on the soundness of the linear PCP

Candidate Linear-Only Encryption from Lattices

[BISW17, GMNO18]

Conjecture: Regev encryption is linear-only

KeyGen(1^λ): Outputs a secret key $\mathbf{s} \in \mathbb{Z}_q^n$

Encrypt($\mathbf{s}, \mu \in \mathbb{Z}_p$): Sample random $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, error $e \leftarrow \chi$ and output
$$\text{ct} = (\mathbf{a}, \mathbf{s}^T \mathbf{a} + pe + \mu)$$

Decrypt(\mathbf{s}, ct): Write $\text{ct} = (\mathbf{a}, b)$ and output
$$(b - \mathbf{s}^T \mathbf{a} \bmod q) \bmod p$$

Correct as long as $|e| \leq \frac{q}{2p}$

Candidate Linear-Only Encryption from Lattices

[BISW17, GMNO18]

Conjecture: Regev encryption is linear-only

KeyGen(1^λ): Outputs a secret key $\mathbf{s} \in \mathbb{Z}_q^n$

Encrypt($\mathbf{s}, \mu \in \mathbb{Z}_p$): Sample random $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, error $e \leftarrow \chi$ and output
$$\text{ct} = (\mathbf{a}, \mathbf{s}^T \mathbf{a} + pe + \mu)$$

Decrypt(\mathbf{s}, ct): Additive homomorphism:

- $\text{ct}_1 = (\mathbf{a}_1, \mathbf{s}^T \mathbf{a}_1 + pe_1 + \mu_1)$
- $\text{ct}_2 = (\mathbf{a}_2, \mathbf{s}^T \mathbf{a}_2 + pe_2 + \mu_2)$

Then:

$$\text{ct}_1 + \text{ct}_2 = (\mathbf{a}_1 + \mathbf{a}_2, \mathbf{s}^T (\mathbf{a}_1 + \mathbf{a}_2) + p(e_1 + e_2) + (\mu_1 + \mu_2))$$

Homomorphic operations increase noise growth

Candidate Linear-Only Encryption from Lattices

[BISW17, GMNO18]

Conjecture: Regev encryption is linear-only

KeyGen(1^λ): Outputs a secret key $\mathbf{s} \in \mathbb{Z}_q^n$

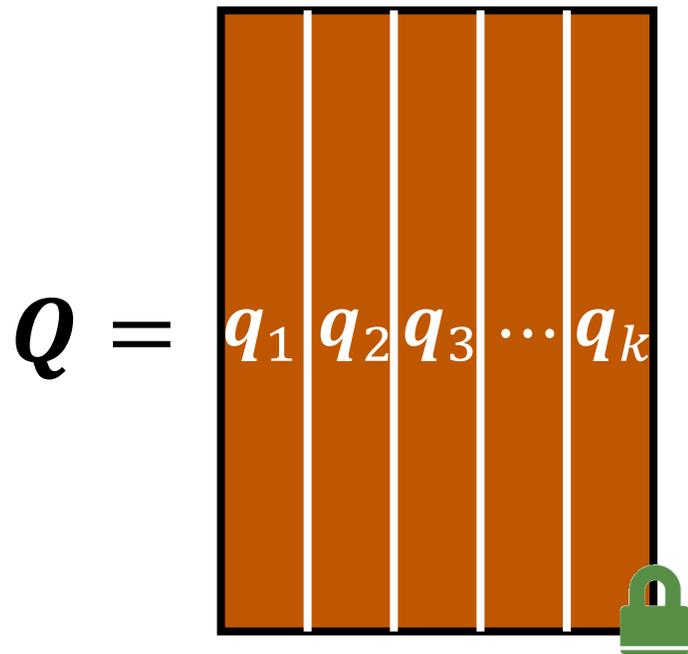
Encrypt($\mathbf{s}, \mu \in \mathbb{Z}_p$): Sample random $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, error $e \leftarrow \chi$ and output
$$\text{ct} = (\mathbf{a}, \mathbf{s}^T \mathbf{a} + pe + \mu)$$

Decrypt(\mathbf{s}, ct): Write $\text{ct} = (\mathbf{a}, b)$ and output
$$(b - \mathbf{s}^T \mathbf{a} \bmod q) \bmod p$$

While Regev encryption can be extended to obtain FHE, existing constructions require additional components or different message embedding

Can we get more homomorphism from vanilla Regev?

Concrete Efficiency of Basic Instantiation



common reference string

homomorphic
evaluation

linear combinations
of length m over \mathbb{F}_p

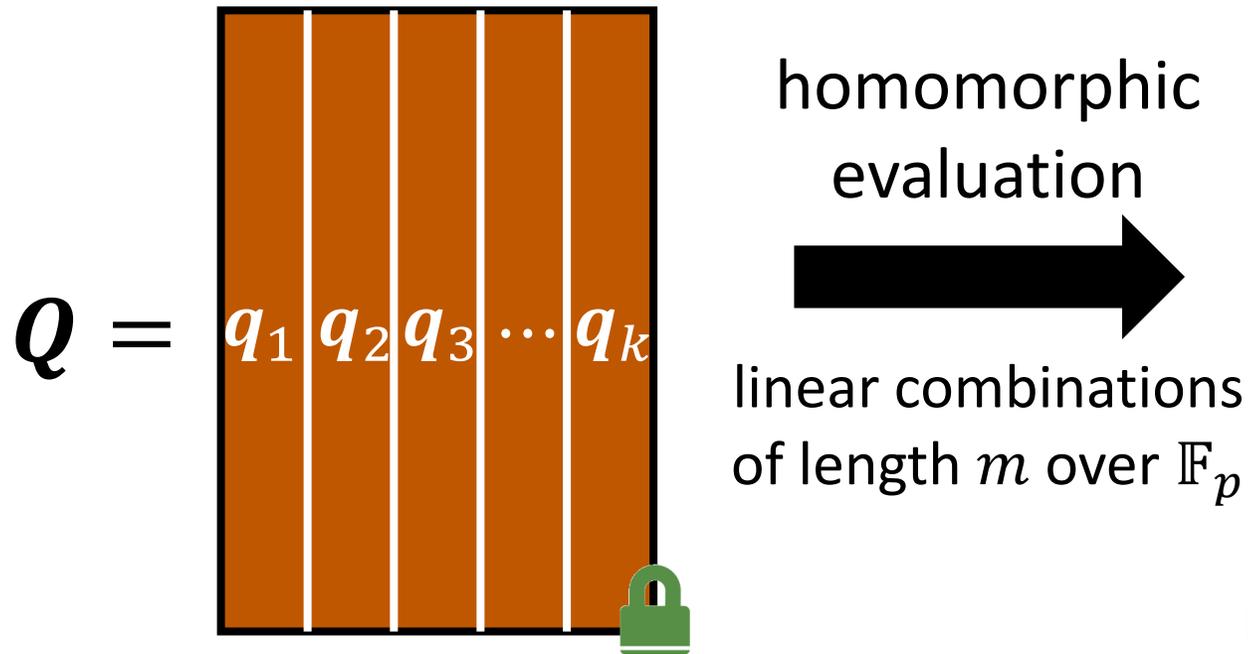
Amount of homomorphism
determines scheme parameters



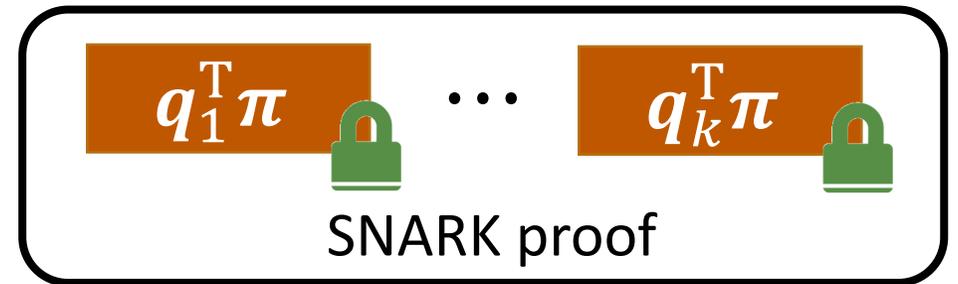
Using quadratic arithmetic programs (for verifying circuit C):

- $k = 4$
- $m = O(|C|)$
- soundness $\approx \frac{2|C|}{|\mathbb{F}_p|} = \frac{2|C|}{p}$

Concrete Efficiency of Basic Instantiation



Amount of homomorphism determines scheme parameters



Need to choose encryption modulus q to support this amount of homomorphism:

$$q/2p > p \cdot m \cdot B$$

where B is the initial noise term

Using quadratic arithmetic programs (for verifying circuit C):

- $k = 4$
- $m = O(|C|)$
- soundness $\approx \frac{2|C|}{|\mathbb{F}_p|} = \frac{2|C|}{p}$

Concrete Efficiency of Basic Instantiation

For a circuit with $m = 2^{20}$ gates and requiring 128 bits of soundness, we require:

- $p > 2^{148}$, so $q > 2^{300}$
- At 128 bits of security, lattice dimension $n > 10^4$, so a single Regev ciphertext is over 350 KB (longer than other post-quantum constructions based on IOPs)
- Proof contains k ciphertexts, so proof is even longer

Alternatively: Use a small plaintext field \mathbb{F}_p and amplify soundness via parallel repetition

- $p \approx 2^{20}$ and $q \approx 2^{100}$: single ciphertext is 45 KB
- Need many copies in this case (≈ 128 copies), so proof is again very long

[GMNO18]: use an instantiation where $p = 2^{32}$ *without* soundness amplification

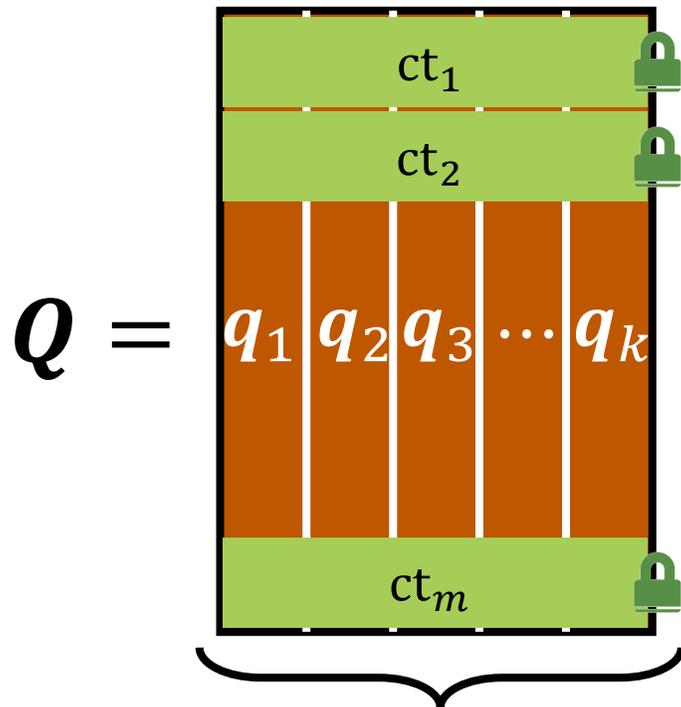
- Proofs are already 640 KB (and provide ≈ 15 bits of provable soundness for verifying computations of size 2^{16})

New techniques needed to reduce proof size

Revisiting the Bitansky et al. Compiler

[BISW17]

Oblivious verifier can “commit”
to its queries ahead of time



part of the CRS



Honest prover takes
(x, w) and constructs
linear PCP $\pi \in \mathbb{F}^m$ and
computes $Q^T \pi$

Key idea: Instead of encrypting
each component of Q individually,
encrypt rows instead

Linear-Only Vector Encryption

[BISW17]

$$v_1 \in \mathbb{F}^k$$

$$v_2 \in \mathbb{F}^k$$

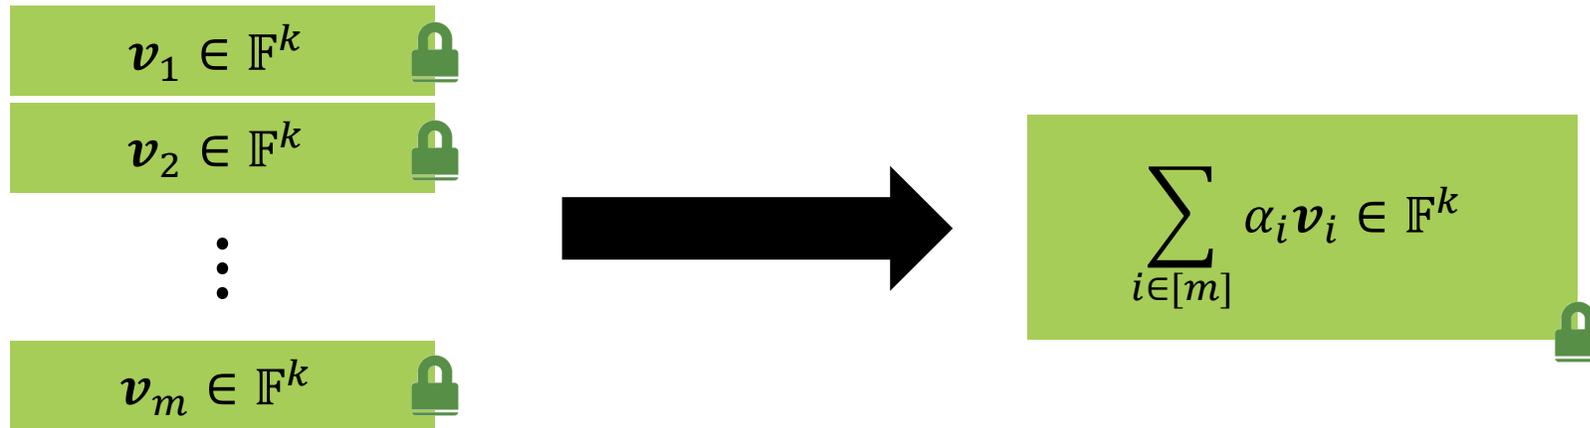
⋮

$$v_m \in \mathbb{F}^k$$

plaintext space is a
vector space

Linear-Only Vector Encryption

[BISW17]



supports homomorphic
vector addition

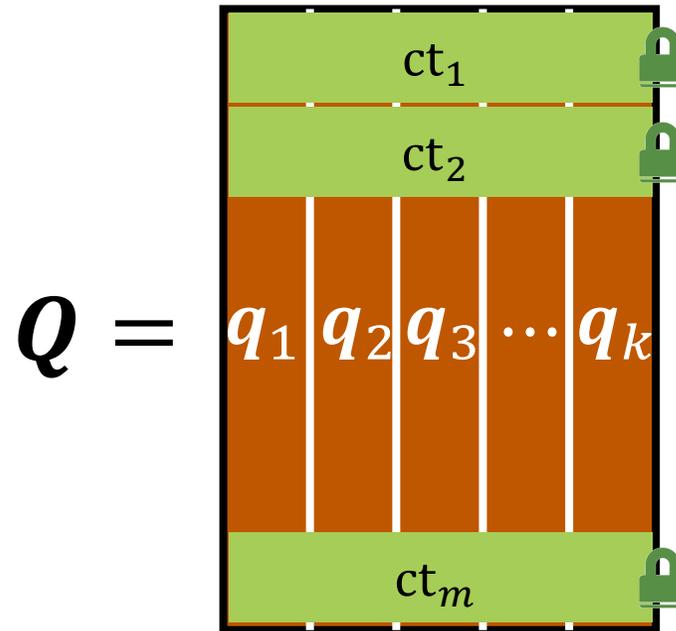
Linear-only: scheme only supports linear homomorphism

From Linear PCPs to Preprocessing SNARGs

[BCIOP13, BISW17]

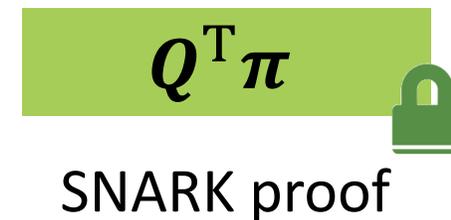


common reference string



Honest prover takes (x, w) and constructs linear PCP $\pi \in \mathbb{F}^m$ and computes $Q^T \pi$

homomorphic evaluation



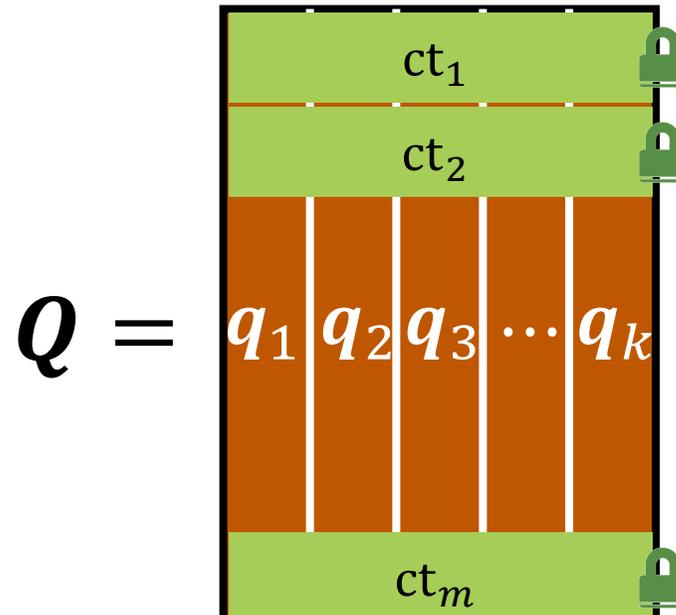
Verifier decrypts to learn $Q^T \pi$ and runs linear PCP decision procedure

From Linear PCPs to Preprocessing SNARGs

[BCIOP13, BISW17]

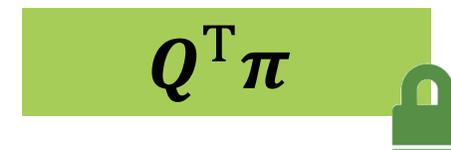


common reference string



Honest prover takes (x, w) and constructs linear PCP $\pi \in \mathbb{F}^m$ and computes $Q^T \pi$

homomorphic evaluation



SNARK proof

- Proof is a single vector encryption ciphertext
- Allows direct compilation from linear PCPs to SNARKs (without extra linearity check from [BCIOP13])

Candidate Linear-Only Vector Encryption

[BISW17]

Conjecture: Regev encryption is linear-only

KeyGen(1^λ): Outputs a secret key $\mathbf{s} \in \mathbb{Z}_q^n$

Encrypt($\mathbf{s}, \mu \in \mathbb{Z}_p$): Sample random $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, error $e \leftarrow \chi$ and output
$$\text{ct} = (\mathbf{a}, \mathbf{s}^T \mathbf{a} + pe + \mu)$$

Decrypt(\mathbf{s}, ct): Write $\text{ct} = (\mathbf{a}, b)$ and output
$$(b - \mathbf{s}^T \mathbf{a} \bmod q) \bmod p$$

Key observation: the same vector $\mathbf{a} \in \mathbb{Z}_q^n$ can be reused with many different secret keys

Amortized/vectorized variant of Regev encryption [PVW08]

Candidate Linear-Only Vector Encryption

[BISW17]

Conjecture: **Vectorized** Regev encryption [PVW08] is linear-only

KeyGen(1^λ): Outputs a secret key $\mathbf{s} \in \mathbb{Z}_q^n$

Encrypt($\mathbf{s}, \mu \in \mathbb{Z}_p$): Sample random $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, error $e \leftarrow \chi$ and output
$$\text{ct} = (\mathbf{a}, \mathbf{s}^T \mathbf{a} + pe + \mu)$$

Decrypt(\mathbf{s}, ct): Write $\text{ct} = (\mathbf{a}, b)$ and output
$$(b - \mathbf{s}^T \mathbf{a} \bmod q) \bmod p$$

Candidate Linear-Only Vector Encryption

[BISW17]

Conjecture: **Vectorized** Regev encryption [PVW08] is linear-only

KeyGen(1^λ): Outputs a secret key $\mathbf{s} \in \mathbb{Z}_q^{n \times k}$

Encrypt($\mathbf{s}, \mu \in \mathbb{Z}_p$): Sample random $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, error $e \leftarrow \chi$ and output
$$\text{ct} = (\mathbf{a}, \mathbf{s}^T \mathbf{a} + pe + \mu)$$

Decrypt(\mathbf{s}, ct): Write $\text{ct} = (\mathbf{a}, b)$ and output
$$(b - \mathbf{s}^T \mathbf{a} \bmod q) \bmod p$$

Candidate Linear-Only Vector Encryption

[BISW17]

Conjecture: **Vectorized** Regev encryption [PVW08] is linear-only

KeyGen(1^λ): Outputs a secret key $\mathbf{s} \in \mathbb{Z}_q^{n \times k}$

Encrypt($\mathbf{s}, \mu \in \mathbb{Z}_p^k$): Sample random $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, error $\mathbf{e} \leftarrow \chi^k$ and output
$$\text{ct} = (\mathbf{a}, \mathbf{s}^T \mathbf{a} + p\mathbf{e} + \mu)$$

Decrypt(\mathbf{s}, ct): Write $\text{ct} = (\mathbf{a}, b)$ and output
$$(b - \mathbf{s}^T \mathbf{a} \bmod q) \bmod p$$

Candidate Linear-Only Vector Encryption

[BISW17]

Conjecture: **Vectorized** Regev encryption [PVW08] is linear-only

KeyGen(1^λ): Outputs a secret key $\mathbf{S} \in \mathbb{Z}_q^{n \times k}$

Encrypt($\mathbf{S}, \mu \in \mathbb{Z}_p^k$): Sample random $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, error $\mathbf{e} \leftarrow \chi^k$ and output
$$\text{ct} = (\mathbf{a}, \mathbf{S}^T \mathbf{a} + p\mathbf{e} + \mu)$$

Decrypt(\mathbf{S}, ct): Write $\text{ct} = (\mathbf{a}, \mathbf{v})$ and output
$$(\mathbf{v} - \mathbf{S}^T \mathbf{a} \bmod q) \bmod p$$

$$|\text{ct}| = (n + k) \log q$$

Would be $k(n + 1) \log q$ using vanilla Regev

Ciphertext size is additive in the vector dimension

Candidate Linear-Only Vector Encryption

[BISW17]

Conjecture: **Vectorized** Regev encryption [PVW08] is linear-only

KeyGen(1^λ): Outputs a secret key $\mathbf{S} \in \mathbb{Z}_q^{n \times k}$

Encrypt($\mathbf{S}, \mu \in \mathbb{Z}_p^k$): Sample random $\mathbf{a} \leftarrow \mathbb{Z}_q^n$, error $\mathbf{e} \leftarrow \chi^k$ and output
$$\text{ct} = (\mathbf{a}, \mathbf{S}^T \mathbf{a} + p\mathbf{e} + \mu)$$

Decrypt(\mathbf{S}, ct): Write $\text{ct} = (\mathbf{a}, \mathbf{v})$ and output
$$(\mathbf{v} - \mathbf{S}^T \mathbf{a} \bmod q) \bmod p$$

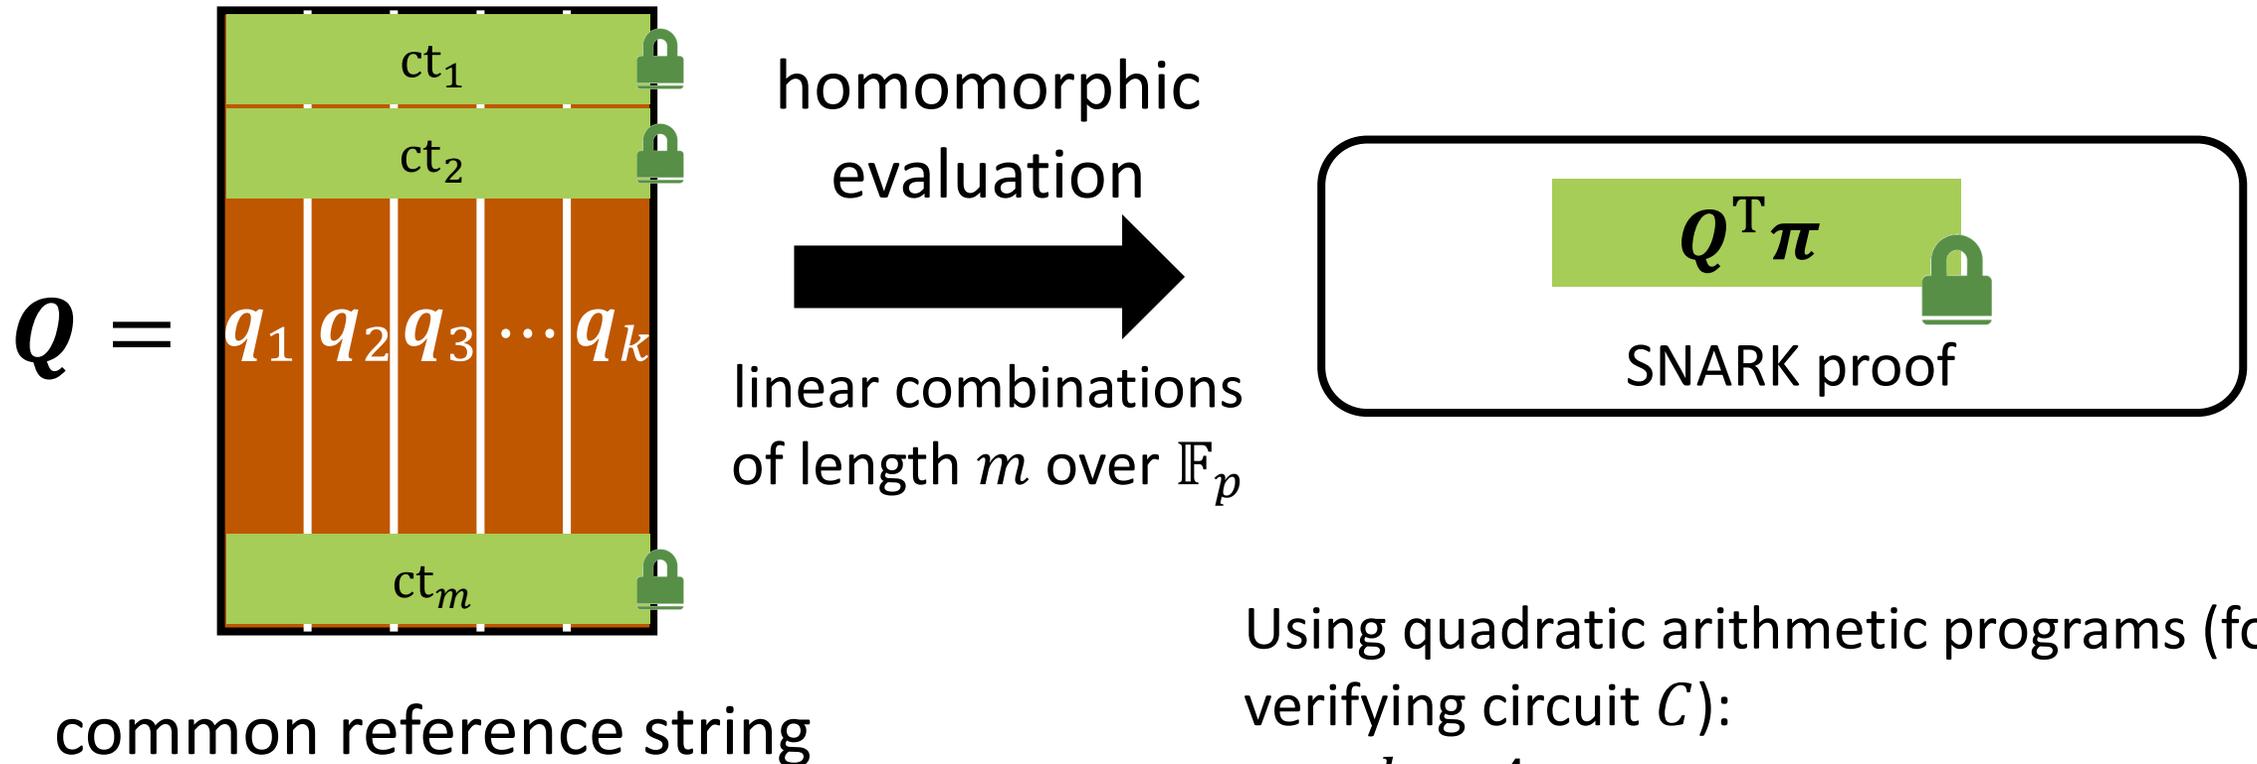
$|\text{ct}| = (n +$

Can use modulus switching [BV11, BGV12] to reduce ciphertext size
after homomorphic evaluation: $(n + k) \log q \rightarrow (n + k) \log q'$

Ciphertext size is additive in the vector dimension

Lattice-Based zkSNARKs using Vector Encryption

[BISW17, ISW21]



Using quadratic arithmetic programs (for verifying circuit C):

- $k = 4$
- $m = O(|C|)$
- soundness $\approx \frac{2|C|}{|\mathbb{F}_p|} = \frac{2|C|}{p}$

Lattice-Based zkSNARKs using Vector Encryption

[BISW17, ISW21]

Previously techniques to achieve small soundness:

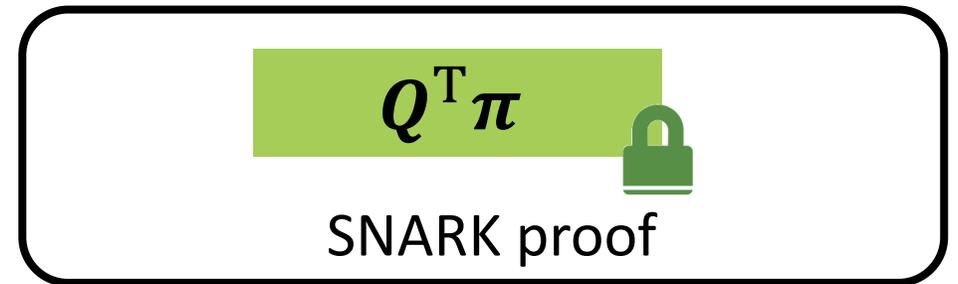
1. Use large p (to ensure LPCP soundness); or
2. Use small p and parallel repetition to amplify soundness

Our approach: parallel repetition of LPCP to amplify soundness:

- Define LPCP to be t independent sets of queries
- Accept only if all t sets accept
- Requires kt LPCP queries and provides soundness $\left(\frac{|C|}{2p}\right)^t$

With vanilla [BCIOP13], same proof size as parallel repetition

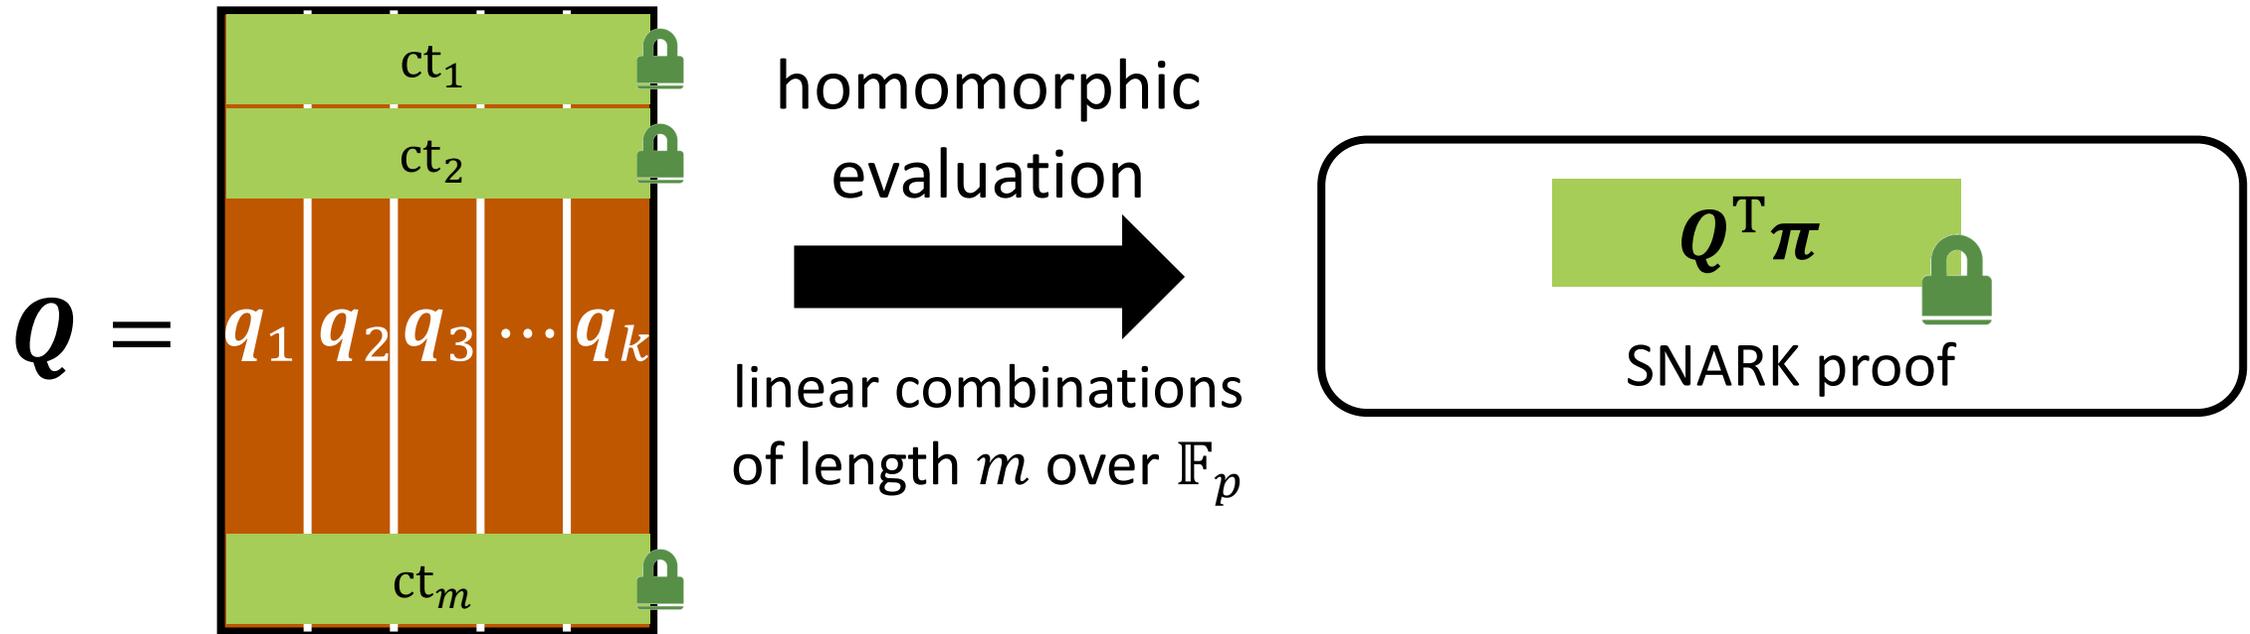
With vector encryption, proof is always a single vector encryption ciphertext and $|ct|$ is *additive* in vector dimension (not multiplicative)



Setting $p \approx 2^{28}$, proof size is 29 KB (with a CRS of size 2.7 GB) for verifying circuit of size 2^{20}

Further Compression via Extensions Fields

[ISW21]



Recall: Noise growth in ciphertexts scales with

- Length m of linear combination
- Magnitude of coefficients in linear combination p

Soundness of linear PCP: $\frac{2|C|}{|\mathbb{F}|}$

Can we further reduce p ?

Idea: use an extension field of small characteristic

Further Compression via Extensions Fields

[ISW21]

(x, w)



$\pi \in \mathbb{F}^m$

linear PCP

Suppose $\mathbb{F} = \mathbb{F}_{p^k}$ where $k > 1$

Can still instantiate using quadratic arithmetic programs

Two approaches to compile to a SNARK:

- Compile LPCP over \mathbb{F}_{p^k} to a LPCP over \mathbb{F}_p , apply linear-only vector encryption over \mathbb{F}_p

Recall that $\mathbb{F}_{p^k} \cong \mathbb{F}_p^k$; field operations in \mathbb{F}_{p^k} are linear transformations over \mathbb{F}_p^k

Transformation increases number of queries and query dimension by k

- Apply linear-only vector encryption over \mathbb{F}_{p^k}

Work over a polynomial ring $R = \mathbb{Z}[x]/\Phi_m(x)$ where m is chosen so that $R/pR \cong \mathbb{F}_{p^k}$

Consider Regev encryption over R (using module lattices)

Further Compression via Extensions Fields

[ISW21]

(x, w)



$\pi \in \mathbb{F}^m$

linear PCP

Suppose $\mathbb{F} = \mathbb{F}_{p^k}$ where $k > 1$

Can still instantiate using quadratic arithmetic programs

Two approaches to compile to a SNARK:

- Compile LPCP over \mathbb{F}_{p^k} to a LPCP over \mathbb{F}_p

Recall that $\mathbb{F}_{p^k} \cong \mathbb{F}_p^k$; field operations in \mathbb{F}_{p^k}

Transformation increases number of queries

- Apply linear-only vector encryption over \mathbb{F}_p

Work over a polynomial ring $R = \mathbb{Z}[x]/\Phi_m$

Consider Regev encryption over R (using mod q)

In both settings: coefficients of prover's linear combination have magnitude $\approx p$ while field has size p^k

Further Compression via Extensions Fields

[ISW21]

(x, w)



$\pi \in \mathbb{F}^m$

linear PCP

This work: consider quadratic extension fields

- $R = \mathbb{Z}[x]/(x^2 + 1)$ and set $p = 3 \pmod{4}$ so $R_p = R/pR \cong \mathbb{F}_{p^2}$
- Choose ciphertext modulus q to be a power of 2
 - All arithmetic operations can be implemented using 128-bit arithmetic
 - Low degree means polynomial arithmetic only slightly more expensive

Further Compression via Extensions Fields

[ISW21]

(x, w)



$\pi \in \mathbb{F}^m$

linear PCP

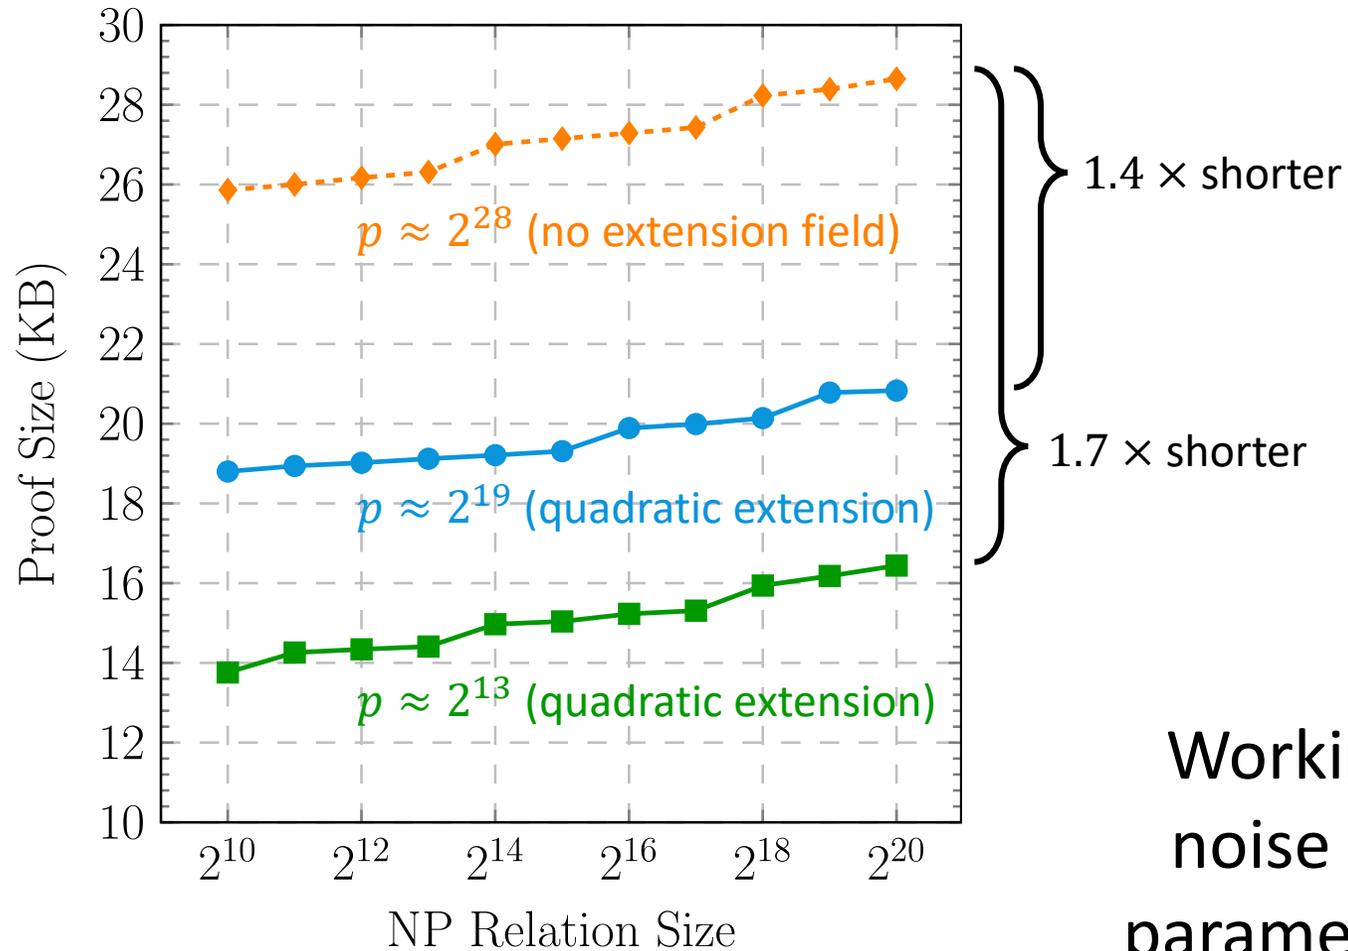
This work: consider quadratic extension fields

- $R = \mathbb{Z}[x]/(x^2 + 1)$ and set $p = 3 \pmod{4}$ so
- Choose ciphertext modulus q to be a power of 2
 - All arithmetic operations can be implemented using 128-bit arithmetic
 - Low degree means polynomial arithmetic only slightly more expensive
- Choose $p = 2^t \pm 1$ so \mathbb{F}_{p^2} has 2^{t+1} -th roots of unity (for efficient implementation of LPCP prover)

Higher-degree extension makes polynomial arithmetic more costly (or need non-power-of-two modulus to exploit FFTs)

Further Compression via Extensions Fields

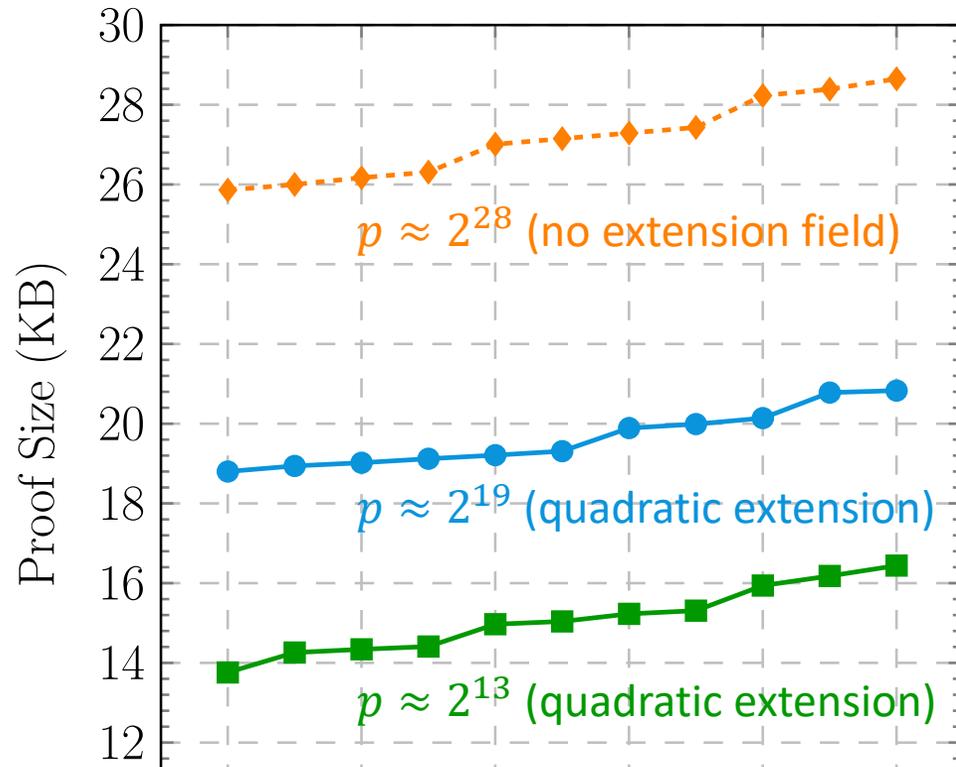
[ISW21]



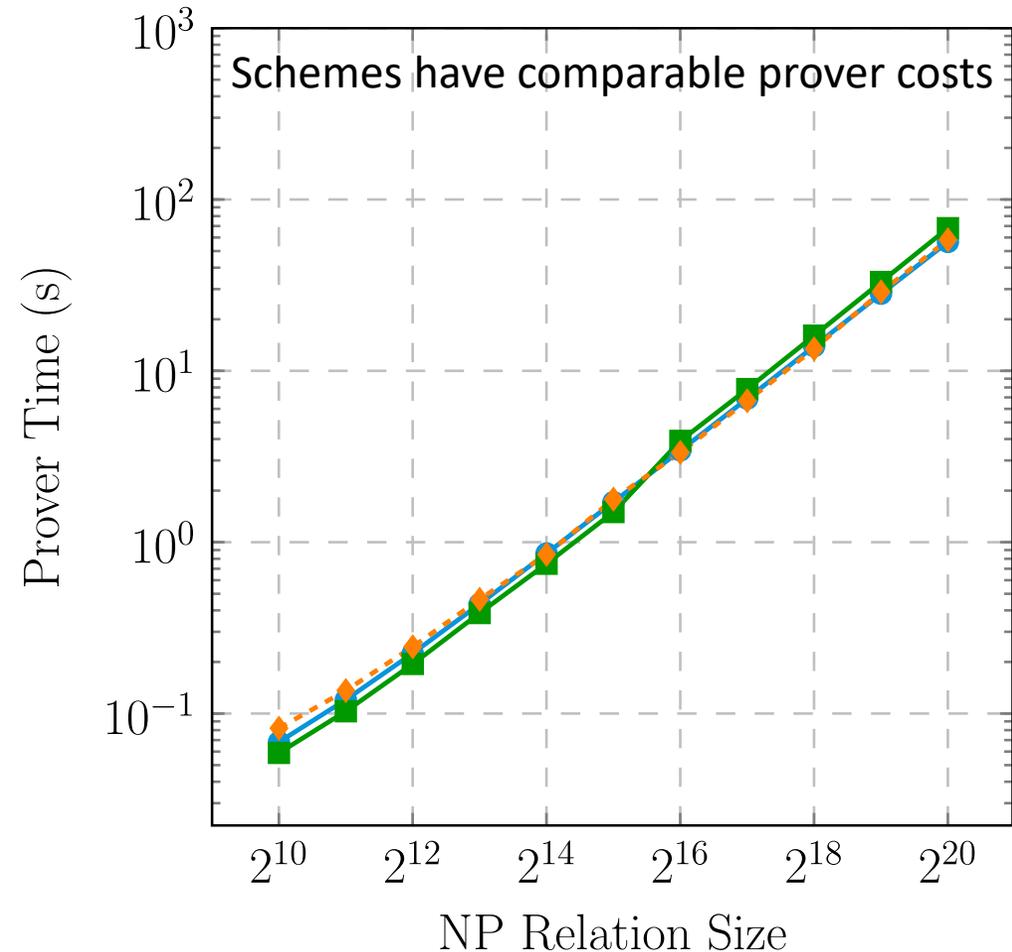
Working over extension field reduces noise accumulation \Rightarrow smaller lattice parameters \Rightarrow concretely shorter proofs

Further Compression via Extensions Fields

[ISW21]

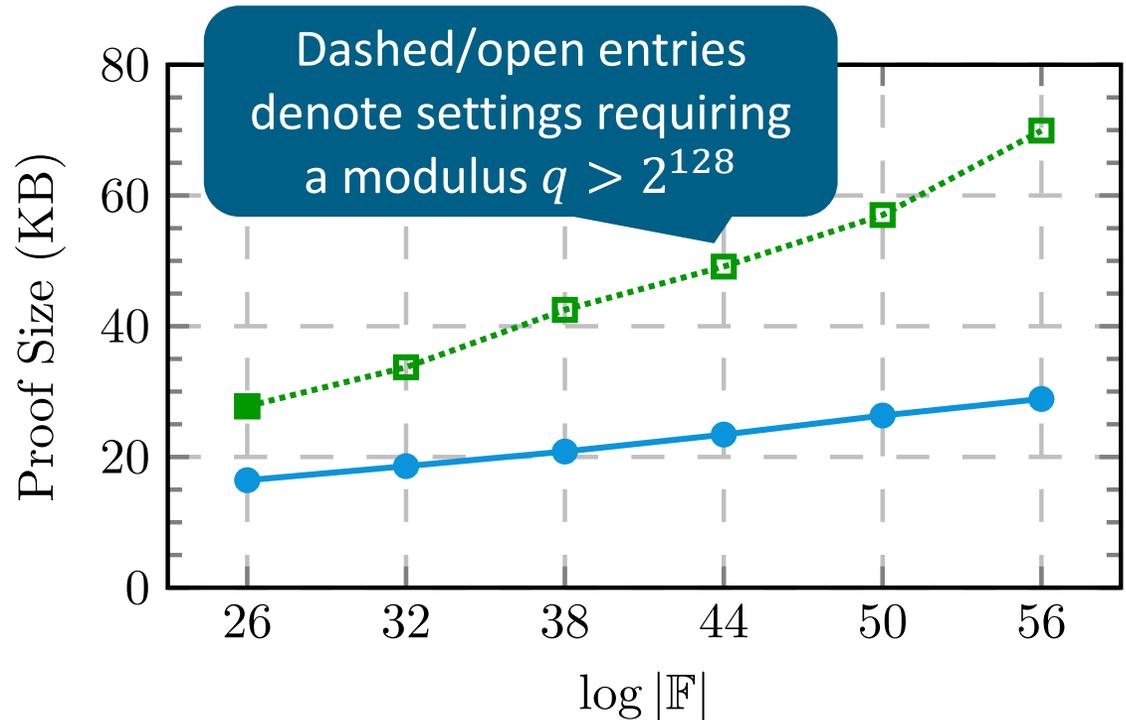
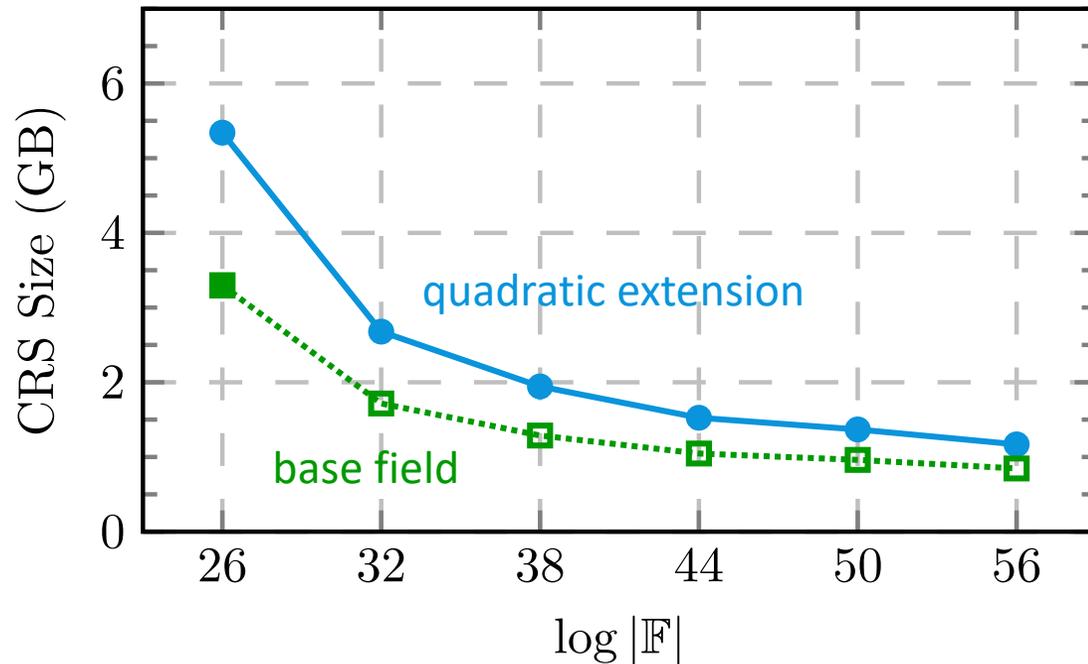


- Slightly more expensive homomorphic operations over extension field, but smaller lattice parameters
- Smaller field \Rightarrow more LPCP queries for soundness amplification \Rightarrow higher prover cost



Effect of Field Size

[ISW21]



Using the extension field increases CRS size but decreases proof size

- CRS consists of “compressed” ciphertexts where random component is derived from a PRF (i.e., $ct = (\mathbf{a}, \mathbf{v})$ where \mathbf{a} is random and $\mathbf{v} = \mathbf{S}^T \mathbf{a} + p\mathbf{e} + \boldsymbol{\mu}$)
- Proof consists of full ciphertexts

[see paper for more microbenchmarks]

Concrete Comparison with zkSNARKs

[ISW21]

Construction	Size		Time			Assumption
	CRS	Proof	Setup	Prover	Verifier	
[Gro16]	199 MB	128 bytes	72 s	79 s	3.4 ms	Pairings <i>Pre-Quantum</i>
Ligero [AHIV17]	–	14 MB	–	38 s	22 s	Random Oracle
Aurora [BCR ⁺ 19]	–	169 KB	–	304 s	6.3 s	Random Oracle
Fractal [COS20]	11 GB	215 KB	116 s	184 s	9.5 ms	Random Oracle
This work	5.3 GB	16.4 KB	2240 s	68 s	1.2 ms	Lattices
This work	1.9 GB	20.8 KB	877 s	56 s	0.4 ms	Lattices <i>Post-Quantum</i>

All benchmarks collected on same hardware for verifying NP relation of size 2^{20}

Concrete Comparison with zkSNARKs

[ISW21]

Construction	Size		Time			Assumption
	CRS	Proof	Setup	Prover	Verifier	
[Gro16]	199 MB	128 bytes				
Ligero [AHIV17]	–	14 MB				
Aurora [BCR ⁺ 19]	–	169 KB				
Fractal [COS20]	11 GB	215 KB				
This work	5.3 GB	16.4 KB				
This work	1.9 GB	20.8 KB	877 s	56 s	0.4 ms	Lattices

Over 10.3× shorter than other post-quantum SNARKs
Still over 131× longer than pairing-based SNARKs
Over 42× shorter than previous lattice-based SNARKs [GMNO18] (based on reported numbers for verifying circuit of size 2^{16})

Post-Quantum

All benchmarks collected on same hardware for verifying NP relation of size 2^{20}

Concrete Comparison with zkSNARKs

[ISW21]

Construction	Size			Time		Assumption
	CRS	Proof	Setup	Prover	Verifier	
[Gro16]	190 MB	190 MB	190 s	79 s	3.4 ms	Pairings Pre-Quantum
Ligero [AHIV17]				38 s	22 s	If we consider restricted computations, can have much faster provers (e.g., ethSTARK [BBHR19])
Aurora [BCR ⁺ 19]		1.2× faster than pairing-based SNARKs		304 s	6.3 s	
Fractal [COS20]				184 s	9.5 ms	
This work		Slower than schemes like Ligero based on MPC-in-the-head (which does not have succinct verification)		68 s	1.2 ms	
This work				56 s	0.4 ms	Lattices Post-Quantum

All benchmarks collected on same hardware for verifying NP relation of size 2^{20}

Concrete Comparison with zkSNARKs

[ISW21]

Construction	Size		Time		Assumption
	CRS	Proof	Setup	Prover	
[Gro16]	198	198	198	3.4 ms	Pairings <i>Pre-Quantum</i>
Ligero [AHIV17]				22 s	Random Oracle
Aurora [BCR ⁺ 19]				6.3 s	Random Oracle
Fractal [COS20]				9.5 ms	Random Oracle
This work				1.2 ms	Lattices
This work				0.4 ms	Lattices <i>Post-Quantum</i>

Lattice-based SNARKs have very lightweight verification: computing a matrix-vector product ($\approx 200,000$ integer multiplications) and rounding

Well-suited for lightweight or energy-constrained devices

All benchmarks collected on same hardware for verifying NP relation of size 2^{20}

Concrete Comparison with zkSNARKs

[ISW21]

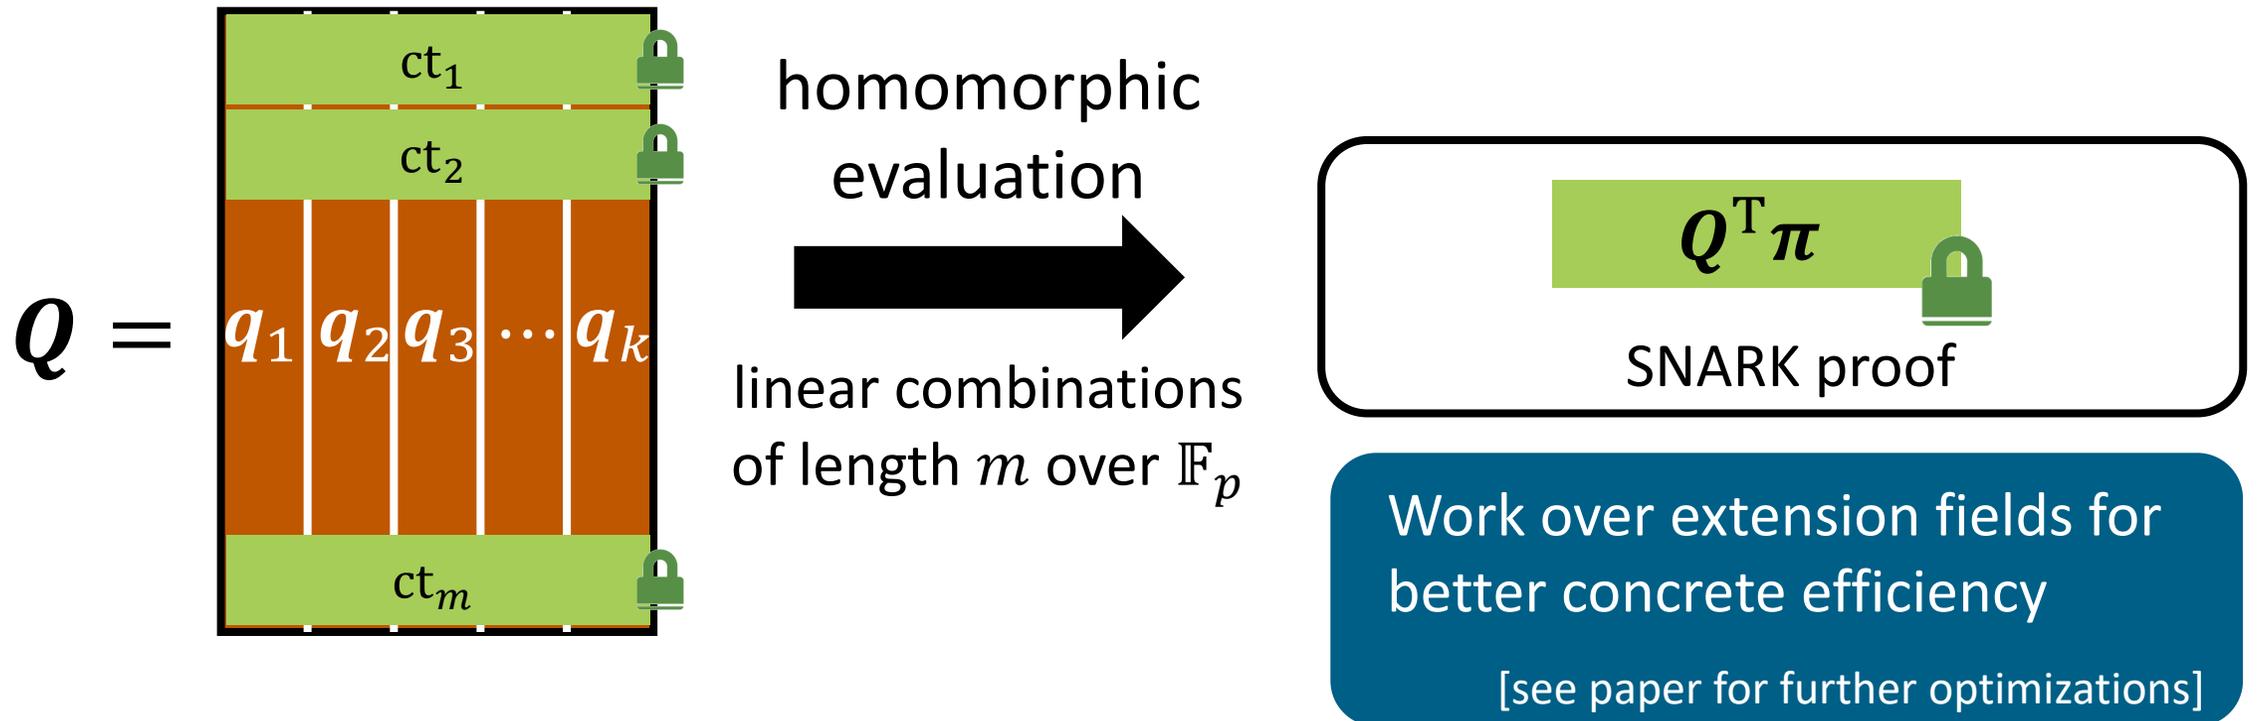
Construction	Size			Time		Assumption
	CRS	Proof	Setup	Prover	Verifier	
[Gro16]	199 MB	128 bytes	72 s	79 s	79 s	Limitations of lattice-based SNARKs: <ul style="list-style-type: none">• Resulting construction is designated-verifier (other schemes are publicly-verifiable)• Require expensive trusted setup (need to encrypt large number of vectors)• Resulting CRS is large (lattice ciphertexts still large, even with compression)
Ligero [AHIV17]	–	14 MB	–	–	–	
Aurora [BCR ⁺ 19]	–	169 KB	–	–	–	
Fractal [COS20]	11 GB	215 KB	116 s	116 s	116 s	
This work	5.3 GB	16.4 KB	2240 s	2240 s	2240 s	
This work	1.9 GB	20.8 KB	877 s	877 s	877 s	

Post-Quantum

All benchmarks collected on same hardware for verifying NP relation of size 2^{20}

Summary

Directly compile linear PCPs to SNARKs using linear-only vector encryption
Instantiate linear-only vector encryption from vectorized Regev encryption



Open Problems

Concretely-efficient **publicly-verifiable** SNARKs from lattices

Constructions with short proofs but expensive verifiers are known from lattices
[BBC⁺18, BLNS20]

Concretely-efficient designated-verifier SNARKs with **reusable soundness** from lattices

Thank you!

<https://eprint.iacr.org/2021/977>

<https://github.com/lattice-based-zkSNARKs/lattice-zksnark>