A Transparent Post-Quantum Zero-Knowledge Argument for R1CS

Vu Hung Tung

School of Information and Communication Technology,
Hanoi University of Science and Technology, Hanoi, Vietnam
s1rjam3s2006@gmail.com

Abstract. We construct a transparent, post-quantum zero-knowledge argument of knowledge for R1CS. Satisfiability of n constraints is reduced by two masked sumchecks to linear claims on a committed short vector; every field element the prover reveals is uniform. The witness is committed row-wise with Module-SIS (Ajtai) commitments; evaluation claims are proven by folding the rows with monomial challenges, whose differences have short scaled inverses, so that the extraction norm grows only singly exponentially in the folding depth. The row commitments and partial evaluations are moved into the witness of a LaBRADOR proof and replaced by two short outer commitments, so the proof has size polylogarithmic in n. The verifier runs in time $\tilde{O}(\sqrt{n}) \cdot \mathrm{poly}(\lambda)$ plus T_{circ}, the time to evaluate the multilinear extensions of the constraint matrices, which we bound for data-parallel circuits. We prove completeness, knowledge soundness of the interactive protocol, and classical zero-knowledge in the programmable random-oracle model, and give parameters from a Core-SVP estimate. The concrete verifier cost at $n = 2^{20}$ exceeds that of checking the circuit directly; the gain is asymptotic. Soundness of the Fiat–Shamir compiled protocol and a quantum-oracle zero-knowledge bound are not proven.

Keywords: post-quantum zero-knowledge · sumcheck · Module-SIS · R1CS · lattice commitments

1 Introduction

1.1 Background

Zero-knowledge proofs [10], introduced by Goldwasser, Micali, and Rackoff, allow a *Prover* to convince a *Verifier* to the validity of a statement without revealing any additional information beyond the truth of the statement itself. The current standard for practical ZK proof system is the ZK-SNARK (*Zero-Knowledge Succinct Non-Interactive Arguement of Knowledge*), in which a proof is succinct when its size and the verifier's work are sublinear in the size of the computation, and non-interactive when it consists of a singular message from Prover to Verifier. The most efficient ZK-SNARKS today - Groth16 [2], PLONK [11], and their

variants - achieve constant-size proofs with constant-time verification. However, they all rely on the hardness of the discrete logarithm problem in the bilinear groups, a problem efficiently solved by Shor's quantum algorithm [12]. Hence, their security is broken in the presence of a sufficiently powerful quantum computer. NIST completed its post-quantum cryptography standardisation process in 2024, selecting lattice-based schemes (Kyber for key encapsulation, Dilithium and FALCON for signatures) as primary standards. The trend has shifted away from systems whose security relies on classical hardness assumptions towards post-quantum approaches.

1.2 Post-Quantum ZK Proofs

The search for post-quantum ZK proofs with competitive efficiency has produced several lines of work, each with significant drawbacks.

Hash-based transparent arguments (STARKs). Scalable transparent arguments of knowledge (STARKs) [3] achieve $O(\log^2 n)$ proof size and $O(\log^2 n)$ verification based solely on collision-resistant hashing. They require no trusted setup. However, STARKs do not achieve succinctness in the strong sense: the $\log^2 n$ factor in both proof size and verifier time is substantial in practice, and the concrete proof sizes (often hundreds of kilobytes to megabytes) are orders of magnitude larger than pairing-based alternatives. Moreover, STARKs do not natively support committed-witness verification without additional machinery.

Lattice-based arguments. Lattice assumptions — Module-LWE and Module-SIS — reduce to worst-case problems on lattices [6, 7] and are believed to be hard for both classical and quantum adversaries. Several lattice-based ZK proof systems exist. Lattice Bulletproofs [5] adapt the Bulletproofs inner-product argument [4] to module lattices, achieving logarithmic-size proofs but with a verifier whose work remains linear in the statement size. BDLOP-style commitments [17] provide efficient lattice-based commitments and openings but do not by themselves yield proofs of R1CS satisfiability. LaBRADOR [21] and Greyhound [22] give much more compact lattice proofs and commitments.

The trusted setup barrier. Many efficient post-quantum arguments that achieve succinct proofs require a structured reference string (SRS) generated by a trusted party whose randomness must be securely discarded. This trusted setup is a systemic vulnerability: a compromise of the toxic waste during setup — by the setup party, by coercion, or by hardware failure — allows the construction of false proofs that pass verification. For high-stakes applications (government identity systems, financial infrastructure, long-lived credential systems) this is often an unacceptable assumption.

1.3 Contributions

1. **Zero-knowledge sumchecks with a uniform transcript.** The R1CS condition is reduced by a row-check and a column-check sumcheck over $\mathbb{F}_q$

to four linear claims. The polynomials are masked by polynomials vanishing on the Boolean hypercube and by zero-sum round masks, so that every message is exactly uniform (Lemmas 3–4); the starting claim of the second sumcheck is made free by an additional uniform constant tied by one linear claim.

2. **Row-masked evaluation.** The witness is committed as a matrix of short rows. A per-row mask p_i multiplied by a polynomial vanishing on the cube makes the partial evaluations g_i uniform, so they can be sent in the clear (Lemma 8); this removes the obstruction that evaluation weights are large field elements and cannot fold short commitments.
3. **Folding with monomial challenges.** Challenges X^a preserve norms exactly and their differences have scaled inverses of ℓ_1-norm $\leq d$ (Lemma 7). Tree extraction yields openings with multiplier 2^{h+1} and norm $(2d+2)^{h+1}X_0$ (Lemma 12); cross-multiplication of multipliers, which would square norms at each level, never occurs.
4. **Recursion to polylogarithmic proofs.** Row commitments and partial evaluations are digit-decomposed, committed by two short plain Ajtai maps, and proven consistent inside one LaBRADOR instance per repetition (§4.3); this is safe without LaBRADOR's zero-knowledge because everything it receives is simulatable (Theorem 4).
5. **Structured circuits and parameters.** For data-parallel R1CS the evaluator costs $O(k + n_0 + \mathrm{nnz}_0)$ (Lemma 2). Parameters for 128-bit quantum Core-SVP are given for all keys (§5).

We prove completeness, knowledge soundness of the interactive protocol (Theorem 3) and zero-knowledge (Theorem 4).

1.4 Related Work

Spartan [14] reduces R1CS to sumchecks; its zero-knowledge relies on discrete-log polynomial commitments, which we replace with Module-SIS and a uniform transcript. LaBRADOR [21] gives compact lattice proofs for R1CS-type relations but is not zero-knowledge; Greyhound [22] gives polynomial commitments with square-root verification. We use LaBRADOR as a compressor and prove that its missing zero-knowledge is harmless here. Lattice Bulletproofs [5] have linear verifiers. Hash-based systems [3, 16] have no algebraic commitment structure and larger concrete proofs. Ling et al. [1] prove statements about a committed symmetric Boolean function; the relation differs.

1.5 Limitations

(i) At $n = 2^{20}$ our verifier performs about 2×10^9 coefficient operations, more than checking the circuit; by $\sqrt{n}$ scaling it wins only for $n \gtrsim 2^{42}$. (ii) Proofs contain $\nu + 1 \approx 33$ LaBRADOR proofs and ≈ 0.6 MB of short commitments; we do not estimate LaBRADOR's size for our relations. (iii) Fiat–Shamir soundness of the compiled protocol and a quantum-oracle zero-knowledge bound are not

proven. (iv) The evaluator is proven only for data-parallel circuits. (v) Soundness relies on a parallel-repetition theorem for tree-special-sound protocols [25], used as a black box, and on LaBRADOR's knowledge soundness [21]. (vi) Parameters are Core-SVP estimates.

Table 1: Assumptions and cited results.

Item	Used for
M-SIS, five keys (§5)	binding, uniqueness of extracted vectors
LaBRADOR, Thm. 5.1 of [21]	exact extraction of final witnesses
Parallel repetition [25]	knowledge error $\kappa_{\mathrm{f}}^{\nu}$ of the row layer
Schwartz–Zippel, leftover hash lemma	sumchecks; hiding of commitments
Programmable random oracle	zero-knowledge

2 Preliminaries

2.1 R1CS

Definition 1 (R1CS Relation). *The relation $\mathcal{R}_{\mathrm{R1CS}}$ is the set of all pairs $((\mathbb{F}, \kappa, n, m, A, B, C, \mathbf{v}), \mathbf{w})$ where $\mathbb{F}$ is a finite field; $\kappa, n, m \in \mathbb{N}$ denote the number of public inputs, total variables, and constraints respectively with $\kappa < n$; $A, B, C \in \mathbb{F}^{m \times n}$; $\mathbf{v} \in \mathbb{F}^{\kappa}$; $\mathbf{w} \in \mathbb{F}^{n-\kappa-1}$; such that $\forall i \in [m]$:*

$$\Big(\sum_{j=0}^{n-1} A_{i,j} \cdot z_j\Big) \cdot \Big(\sum_{j=0}^{n-1} B_{i,j} \cdot z_j\Big) = \sum_{j=0}^{n-1} C_{i,j} \cdot z_j$$

where $\mathbf{z} \coloneqq (1,\, \mathbf{v},\, \mathbf{w}) \in \mathbb{F}^n$.

Square R1CS assumption. Throughout this work we assume $m = n = 2^k$ for some $k \in \mathbb{N}$, so that a single k-bit boolean hypercube uniformly addresses both constraint rows and witness columns. This is standard in the sumcheck-based SNARK literature and entails no loss of generality for circuits padding to the next power of two.

2.2 Domains

Definition 2 (Power-of-two Cyclotomic Ring). *Let $d = 2^{\ell}$ for some $\ell \geq 1$. The $2d$-th cyclotomic polynomial $\Phi_{2d}(X) = X^d + 1$ is monic and irreducible over $\mathbb{Z}$. The* power-of-two cyclotomic ring *is:*

$$R \;=\; \mathbb{Z}[X]/(X^d + 1).$$

Modulus. We take q prime with $q \equiv 5 \pmod 8$ and $d \geq 2$ a power of two. Then $X^d + 1$ factors over $\mathbb{Z}_q$ into two irreducible factors of degree $d/2$ (of the form $X^{d/2} \pm r$ with $r^2 = -1$), and by [18] every nonzero $y \in R_q$ with $\|y\|_\infty < \sqrt{q/2}$ is invertible. This is needed for extraction (differences of challenges must be invertible) and for the hiding lemma. Multiplication is still fast: a partial NTT with $d/2$-point residues costs $O(d \log d)$. (The fully split choice $q \equiv 1 \pmod{2d}$ would make some small differences non-invertible and is *not* used.) We write $R_q \coloneqq \mathbb{Z}_q[X]/(X^d + 1)$ and $\mathbb{F}_q = \mathbb{Z}_q$.

2.3 Number Theoretic Transform

The NTT serves as the finite-field analogue of the DFT, converting polynomials from coefficient to evaluation representation for efficient arithmetic.

Definition 3 (Primitive Root of Unity). *A scalar $\psi \in \mathbb{Z}_q$ is a* primitive $2d$-th root of unity *if $\psi^{2d} \equiv 1 \pmod q$ and $\psi^j \not\equiv 1 \pmod q$ for all $j \in \{1, \ldots, 2d-1\}$.*

Evaluation set. The d distinct roots of $X^d + 1$ in $\mathbb{Z}_q$ are the odd powers of ψ: $\{\psi^1, \psi^3, \psi^5, \ldots, \psi^{2d-1}\}$.

Definition 4 (NTT and INTT). *For $a(X) = \sum_{i=0}^{d-1} a_i X^i \in R_q$, the* Number Theoretic Transform $\mathrm{NTT} : R_q \to \mathbb{Z}_q^d$ *maps the coefficient vector to its evaluations at the evaluation set:*

$$\hat{a}_m \;\coloneqq\; a(\psi^{2m+1}) \;=\; \sum_{i=0}^{d-1} a_i \, \psi^{(2m+1)i} \pmod q, \quad m = 0, \ldots, d-1.$$

The Inverse NTT $\mathrm{INTT} : \mathbb{Z}_q^d \to R_q$ *recovers coefficients via:*

$$a_i \;\coloneqq\; d^{-1} \sum_{m=0}^{d-1} \hat{a}_m \, \psi^{-(2m+1)i} \pmod q.$$

Convolution theorem and complexity. Polynomial multiplication modulo $X^d + 1$ maps to an entry-wise product in the evaluation domain:

$$\mathrm{NTT}\big(a(X) \cdot b(X) \bmod X^d + 1\big) \;=\; \mathrm{NTT}(a) \odot \mathrm{NTT}(b).$$

Using the Cooley-Tukey butterfly structure, both NTT and INTT run in $O(d \log d)$ operations over $\mathbb{Z}_q$. All ring multiplications in the protocol are computed via NTT for this efficiency gain.

2.4 Metric Norms

Coefficients $a_i \in \mathbb{Z}_q$ are implicitly lifted to the symmetric representative $a_i \in \left[-\frac{q}{2}, \frac{q}{2}\right) \cap \mathbb{Z}$.

Definition 5 (Infinity Norm). *For* $a(X) = \sum_{i=0}^{d-1} a_i X^i \in R_q$*:* $\|a\|_\infty \coloneqq \max_{0 \le i < d} |a_i|$*. For a module vector* $\mathbf{a} = (a_0, \ldots, a_{s-1}) \in R_q^s$*:* $\|\mathbf{a}\|_\infty \coloneqq \max_{0 \le j < s} \|a_j\|_\infty$*.*

Definition 6 (ℓ_1 Norm). *For* $a(X) \in R_q$*:* $\|a\|_1 \coloneqq \sum_{i=0}^{d-1} |a_i|$*. For* $\mathbf{a} \in R_q^s$*:* $\|\mathbf{a}\|_1 \coloneqq \sum_{j=0}^{s-1} \|a_j\|_1$*.*

2.5 Bounded Challenge Space

Definition 7 (Challenge space). $\mathcal{C}_R \coloneqq \{c \in R_q : \|c\|_\infty \le 2\}$*, so* $|\mathcal{C}_R| = 5^d$ *and* $\|c\|_1 \le \omega \coloneqq 2d$*. We require* $5^d \ge 2^\lambda$*, e.g.* $d = 64$ *gives* $|\mathcal{C}_R| \approx 2^{148}$*.*

Differences $c - c'$ of challenges have ℓ_∞-norm $\le 4 < \sqrt{q/2}$, hence are invertible when nonzero (§2.2).

Lemma 1 (Bounded Noise Growth). *For* $e(X) \in R_q$ *and* $c(X) \in \mathcal{C}_R$*, the anti-cyclic convolution in* R_q *satisfies:*

$$\|c \cdot e\|_\infty \;\le\; \|c\|_1 \cdot \|e\|_\infty \;\le\; \omega \cdot \|e\|_\infty.$$

Proof. In R_q with $X^d \equiv -1$, the coefficient of X^j in $c \cdot e$ is $\sum_i c_i e_{j-i \bmod 2d} \cdot (\pm 1)$. Taking absolute values and using $\|c\|_\infty = 1$ with at most ω nonzero terms: $|(c \cdot e)_j| \le \omega \cdot \|e\|_\infty$.

2.6 Embedding scalars as constants

A scalar $x \in \mathbb{Z}_q$ is identified with the constant polynomial $x \in R_q$. A vector $\mathbf{x} \in \mathbb{Z}_q^N$ is thereby a vector in R_q^N with $\|\mathbf{x}\|_\infty$ unchanged. For a public $\mathbf{t} \in \mathbb{Z}_q^N$ the map $\tilde{\mathbf{x}} \mapsto \sum_s t_s \tilde{x}_s$ is R_q-linear, and on constant vectors takes values in $\mathbb{Z}_q \subset R_q$. We call $\mathrm{ct}(f)$ the constant coefficient of $f \in R_q$.

2.7 Multilinear Extensions

Definition 8 (Multilinear Extension (MLE)). *For a function* $f : \{0,1\}^k \to \mathbb{Z}_q$*, its* multilinear extension $\hat{f} : \mathbb{Z}_q^k \to \mathbb{Z}_q$ *is the unique polynomial of degree at most one in each variable satisfying* $\hat{f}(\mathbf{b}) = f(\mathbf{b})$ *for all* $\mathbf{b} \in \{0,1\}^k$*:*

$$\hat{f}(\mathbf{x}) \;\coloneqq\; \sum_{\mathbf{b} \in \{0,1\}^k} f(\mathbf{b}) \cdot \prod_{j=0}^{k-1} \big(x_j b_j + (1 - x_j)(1 - b_j)\big).$$

For an R1CS matrix $M \in \mathbb{Z}_q^{m \times n}$ *with* $m = n = 2^k$*, its 2D MLE* $\hat{M}(\mathbf{y}, \mathbf{x}) \in \mathbb{Z}_q[\mathbf{y}, \mathbf{x}]$ *satisfies* $\hat{M}(\mathbf{b}_{\mathrm{row}}, \mathbf{b}_{\mathrm{col}}) = M_{\mathrm{addr}(\mathbf{b}_{\mathrm{row}}), \mathrm{addr}(\mathbf{b}_{\mathrm{col}})}$ *for all boolean inputs, where* $\mathrm{addr}(\mathbf{b}) = \sum_j b_j 2^j$*.*

Definition 9 (Equality Polynomial).

$$\mathrm{eq}(\boldsymbol{\tau}, \mathbf{x}) \;\coloneqq\; \prod_{j=0}^{k-1} \big(\tau_j x_j + (1 - \tau_j)(1 - x_j)\big).$$

This is the MLE of the Kronecker delta: $\mathrm{eq}(\boldsymbol{\tau}, \mathbf{b}) = 1$ *iff* $\mathbf{b} = \boldsymbol{\tau}$ *on* $\{0,1\}^k$*.*

2.8 Module Lattice Commitment

Definition 10 (Ajtai-type commitment with explicit randomness). *Let* $m, l, l_r \in \mathbb{N}$. $\mathrm{Setup}(1^\lambda)$ *samples* $\mathbf{A} = [\mathbf{A}_1 \mid \mathbf{A}_2] \leftarrow R_q^{m\times(l+l_r)}$ *uniformly (e.g. by expanding a public hash seed; no trapdoor exists). To commit to* $\tilde{\mathbf{u}} \in R_q^l$, *sample* $\mathbf{r} \leftarrow \{-1, 0, 1\}^{l_r d}$ *(as a vector in* $R_q^{l_r}$*) and output* $\mathbf{C} = \mathbf{A}_1\tilde{\mathbf{u}} + \mathbf{A}_2\mathbf{r} \in R_q^m$.

Definition 11 (Relaxed opening and binding). *A* relaxed opening *of* $\mathbf{C}$ *with bound* B *is* $(\bar{\mathbf{s}}, \bar{\mathbf{s}}_r, \bar{c})$ *with* $\bar{c} \neq 0$, $\|\bar{c}\|_\infty \le 4$, $\|\bar{c}\|_1 \le 2\omega$, $\|\bar{\mathbf{s}}\|_\infty, \|\bar{\mathbf{s}}_r\|_\infty \le B$ *and* $\mathbf{A}_1\bar{\mathbf{s}} + \mathbf{A}_2\bar{\mathbf{s}}_r = \bar{c}\,\mathbf{C}$. *The scheme is* relaxed-binding *under M-SIS*$_{q,m,l+l_r,4\omega B}$*: two relaxed openings* $(\bar{\mathbf{s}}^i, \bar{\mathbf{s}}^i_r, \bar{c}^i)$, $i = 1, 2$, *of the same* $\mathbf{C}$ *satisfy* $\bar{c}^2\bar{\mathbf{s}}^1 = \bar{c}^1\bar{\mathbf{s}}^2$ *unless one finds a nonzero M-SIS solution* $\big(\bar{c}^2\bar{\mathbf{s}}^1 - \bar{c}^1\bar{\mathbf{s}}^2,\ \bar{c}^2\bar{\mathbf{s}}^1_r - \bar{c}^1\bar{\mathbf{s}}^2_r\big)$ *of* ℓ_∞*-norm at most* $4\omega B$ *for* $[\mathbf{A}_1|\mathbf{A}_2]$.

Remark. A commitment $\mathbf{A}\tilde{\mathbf{u}}+\mathbf{e}$ with the message as LWE secret is not hiding against distinguishers that know candidate messages; the explicit randomness $\mathbf{r}$ gives statistical hiding (Lemma 5) without an M-LWE assumption.

2.9 Monomial challenges and gadget

Let $\mathcal{S} = \{X^a : 0 \le a < 2d\} \subset R_q$ ($|\mathcal{S}| = 2d$, $X^d = -1$). For a base 2^b and $\tau = \lceil \log_2 q/b \rceil$, $\mathbf{G}^{-1}(\mathbf{C}) \in R_q^{m\tau}$ denotes the coefficientwise balanced base-2^b digits of $\mathbf{C} \in R_q^m$ (entries $\le 2^{b-1}$ in absolute value), and $\mathbf{G} = (1, 2^b, \ldots, 2^{b(\tau-1)}) \otimes I_m$, so $\mathbf{G}\mathbf{G}^{-1}(\mathbf{C}) = \mathbf{C}$. The same notation is used for scalars ($m = 1$).

3 Construction

3.1 Relation and structured circuits

Let $n = 2^k$ and $A, B, C \in \mathbb{F}_q^{n\times n}$. A public index set $P \subset \{0,1\}^k$ contains position $\mathbf{0}$ ($z_{\mathbf{0}} = 1$); the public values z_P are known to the Verifier. The witness is $\mathbf{w} = (z_{\mathbf{j}})_{\mathbf{j}\notin P}$, and the relation is $(A\mathbf{z}) \circ (B\mathbf{z}) = C\mathbf{z}$. For completeness $\|\mathbf{w}\|_\infty \le B_w$; soundness holds for the relation over $\mathbb{F}_q$ without a shortness condition. For $X \in \{A, B, C\}$ write $\tilde{X}(\mathbf{x}) = \widetilde{X\mathbf{z}}(\mathbf{x})$ and $\hat{X}(\mathbf{x}, \mathbf{y}) = \sum_{\mathbf{i},\mathbf{j}} \mathrm{eq}(\mathbf{x}, \mathbf{i})\,\mathrm{eq}(\mathbf{y}, \mathbf{j}) X_{\mathbf{ij}}$, so $\tilde{X}(\mathbf{x}) = \sum_{\mathbf{y}} \hat{X}(\mathbf{x}, \mathbf{y})\hat{z}(\mathbf{y})$. The Verifier needs $M(\mathbf{r}_y) = \sum_X \eta_X \hat{X}(\mathbf{r}_x, \mathbf{r}_y)$ and $\mathrm{pub}(\mathbf{r}_y) = \sum_{\mathbf{j}\in P} \mathrm{eq}(\mathbf{r}_y, \mathbf{j}) z_{\mathbf{j}}$; let T_{circ} be the cost of the former. In general $T_{\mathrm{circ}} = O(\mathrm{nnz})$.

Definition 12 (Data-parallel R1CS). *Let* $k = k_b + k_0$, $B = 2^{k_b}$, $n_0 = 2^{k_0}$, *and* $A_0, B_0, C_0 \in \mathbb{F}_q^{n_0\times n_0}$ *with* nnz_0 *nonzeros in total. The instance is* $A = I_B \otimes A_0$, $B = I_B \otimes B_0$, $C = I_B \otimes C_0$, *with* P *consisting of position* 0 *of each block and the next* κ_0 *positions.*

Lemma 2 (Evaluator). *For* $\mathbf{x} = (\mathbf{x}_b, \mathbf{x}_0)$, $\mathbf{y} = (\mathbf{y}_b, \mathbf{y}_0)$*:* $\hat{A}(\mathbf{x}, \mathbf{y}) = \mathrm{eq}(\mathbf{x}_b, \mathbf{y}_b)\hat{A}_0(\mathbf{x}_0, \mathbf{y}_0)$, *likewise for* B, C. *Hence* $M(\mathbf{r}_y)$ *is computed in* $O(k + n_0 + \mathrm{nnz}_0)$ *operations, independent of* B*; and* $\mathrm{pub}(\mathbf{r}_y)$ *in* $O(k_0 + B\kappa_0 k)$.

Proof. Index $\mathbf{i} = (b, i_0), \mathbf{j} = (b', j_0)$; $A_{\mathbf{ij}} = \delta_{bb'}(A_0)_{i_0 j_0}$ and $\mathrm{eq}(\mathbf{x}, \mathbf{i}) = \mathrm{eq}(\mathbf{x}_b, b)\,\mathrm{eq}(\mathbf{x}_0, i_0)$. So $\hat{A} = \big(\sum_b \mathrm{eq}(\mathbf{x}_b, b)\,\mathrm{eq}(\mathbf{y}_b, b)\big)\hat{A}_0(\mathbf{x}_0, \mathbf{y}_0)$. The first factor is multilinear in each variable and equals $\mathrm{eq}(\mathbf{x}_b, \mathbf{y}_b)$ on the cube, hence everywhere. Tabulating $\mathrm{eq}(\mathbf{x}_0, \cdot), \mathrm{eq}(\mathbf{y}_0, \cdot)$ costs $O(n_0)$; summing over nonzeros costs $O(\mathrm{nnz}_0)$. For pub: $\sum_b \mathrm{eq}(\mathbf{r}_b, b) = 1$, so the constant-1 entries contribute $\mathrm{eq}(\mathbf{r}_{y,0}, \mathbf{0})$; each true public input costs $O(k)$.

For a circuit not of this form take $B = 1$; $T_{\mathrm{circ}} = O(\mathrm{nnz})$ and the verifier is linear. No evaluator is proven for other classes.

3.2 Masks and committed data

Let $\theta^1(x) = 2x - 1$, $\theta^b(x) = x^{b-1}(1-x)$ $(b \geq 2)$. Each satisfies $\theta^b(0) + \theta^b(1) = 0$, and $\theta^1, \ldots, \theta^D$ span $\mathcal{D}_D = \{h \in \mathbb{F}_q[T]_{\leq D} : h(0) + h(1) = 0\}$. The Prover samples uniform scalars α_i^X $(X \in \{A, B, C\},\ i < k)$, $e_{j,b}$ $(j < k,\ b \leq 5)$, σ_0, $e'_{j,b}$ $(j < k,\ b \leq 3)$, in total $11k + 1$, and p_i $(i < \rho)$. With $\rho = 2^h$, $\gamma = n/\rho$, $\mathbf{y} = (\mathbf{y}_h, \mathbf{y}_l)$, $\kappa(\mathbf{y}_l) = y_{l,0}(1 - y_{l,0})$ define

$$\tilde{X}_{\mathrm{zk}}(\mathbf{x}) = \tilde{X}(\mathbf{x}) + \textstyle\sum_{i<k} \alpha_i^X x_i(1 - x_i), \qquad \hat{\rho}_1(\mathbf{x}) = \sum_{j<k, b\leq 5} e_{j,b}\theta^b(x_j),$$
$$\hat{z}_{\mathrm{zk}}(\mathbf{y}) = \hat{z}(\mathbf{y}) + \kappa(\mathbf{y}_l)\hat{p}(\mathbf{y}_h), \qquad \hat{\rho}_2(\mathbf{y}) = \sigma_0 + \textstyle\sum_{j<k, b\leq 3} e'_{j,b}\theta^b(y_j),$$

with $\hat{p}(\mathbf{y}_h) = \sum_i \mathrm{eq}(\mathbf{y}_h, i)p_i$. On $\{0,1\}^k$: $\tilde{X}_{\mathrm{zk}} = \tilde{X}$, $\hat{z}_{\mathrm{zk}} = \hat{z}$, $\sum \hat{\rho}_1 = 0$, $\sum \hat{\rho}_2 = 2^k\sigma_0$.

Committed data. Let $\delta = \lceil \log_2 q \rceil + \lambda$. Each scalar s among the $11k + 1$ mask scalars and each p_i is written as $\sum_{t<\delta} 2^t d_t \bmod q$ with uniform bits (uniform up to $2^{-\lambda}$). *Mask vector:* $\mathbf{u}_M \in \{0,1\}^{N_M}$, $N_M = (11k + 1)\delta$, committed as $\mathbf{C}_M = \mathbf{A}_1^M \mathbf{u}_M + \mathbf{A}_2^M \mathbf{r}_M$ with $\mathbf{r}_M \leftarrow \{-1, 0, 1\}^{l_r^M d}$. *Rows:* arrange the witness (public positions set to 0) as rows $W_i \in \mathbb{F}_q^\gamma$ $(i < \rho)$, set $\mathbf{U}_i = (W_i \| \text{digits of } p_i) \in \mathbb{F}_q^{\gamma'}$, $\gamma' = \gamma + \delta$, and commit $\mathbf{C}_i = \mathbf{A}_1\mathbf{U}_i + \mathbf{A}_2\mathbf{R}_i$, $\mathbf{R}_i \leftarrow \{-1, 0, 1\}^{l_r d}$. All entries are embedded as constants of R_q. Let $\hat{\mathbf{C}}_i = \mathbf{G}^{-1}(\mathbf{C}_i)$ and $\mathbf{T} = \mathbf{B}(\hat{\mathbf{C}}_0 \| \cdots \| \hat{\mathbf{C}}_{\rho-1})$, where $\mathbf{B}$ is a public plain Ajtai matrix. Keys are uniform and expanded from public seeds ($l_r = \lceil (md \log_2 q + 2\lambda)/(d \log_2 3) \rceil$ for each commitment, Lemma 5).

3.3 Protocol

Step 1 (commit). $\mathcal{P} \to \mathcal{V}$: $\mathbf{T}$, $\mathbf{C}_M$.

Step 2 (row check). $\mathcal{V} \to \mathcal{P}$: $\boldsymbol{\tau} \leftarrow \mathbb{F}_q^k$. Let $G_1 = \mathrm{eq}(\boldsymbol{\tau}, \mathbf{x})(\tilde{A}_{\mathrm{zk}}\tilde{B}_{\mathrm{zk}} - \tilde{C}_{\mathrm{zk}}) + \hat{\rho}_1$ (degree ≤ 5 per variable), $S_0 = 0$. For $j < k$: $\mathcal{P}$ sends $\tilde{g}_j(T) = \sum_{\mathbf{b}} G_1(r_{<j}, T, \mathbf{b})$; $\mathcal{V}$ checks $\tilde{g}_j(0) + \tilde{g}_j(1) = S_j$, draws r_j, sets $S_{j+1} = \tilde{g}_j(r_j)$. With $\mathbf{r}_x = (r_j)$, $\mathcal{P}$ sends $V_X = \tilde{X}_{\mathrm{zk}}(\mathbf{r}_x)$; $\mathcal{V}$ sets $a_{\rho 1} = S_k - \mathrm{eq}(\boldsymbol{\tau}, \mathbf{r}_x)(V_A V_B - V_C)$.

Step 3 (column check). $\mathcal{V} \to \mathcal{P}$: $\boldsymbol{\eta} \in \mathbb{F}_q^3$; $M(\mathbf{y}) = \sum_X \eta_X \hat{X}(\mathbf{r}_x, \mathbf{y})$. $\mathcal{P}$ sends $\hat{S}' = \sum_{\mathbf{y}}(M\hat{z}_{\mathrm{zk}} + \hat{\rho}_2)(\mathbf{y}) = \sum_X \eta_X \tilde{X}(\mathbf{r}_x) + 2^k\sigma_0$, sets $S'_0 = \hat{S}'$, and runs the same round protocol on $G_2 = M\hat{z}_{\mathrm{zk}} + \hat{\rho}_2$ (degree ≤ 3), giving $\mathbf{r}_y = (\mathbf{r}_h, \mathbf{r}_l)$. $\mathcal{P}$ sends $a = \hat{z}_{\mathrm{zk}}(\mathbf{r}_y)$, and $\mathbf{T}_0 = \mathbf{B}_0(\mathbf{G}^{-1}(g_0) \| \cdots \| \mathbf{G}^{-1}(g_{\rho-1}))$ where $g_i = \sum_j \mathrm{eq}(\mathbf{r}_l, j)W_{ij} + \kappa(\mathbf{r}_l)p_i$. $\mathcal{V}$ computes $M(\mathbf{r}_y)$ and sets $a_{\rho 2} = S'_k - M(\mathbf{r}_y)a$, $a_T = \hat{S}' - \sum_X \eta_X V_X$.

Mask claims. Let $\mathbf{t}_{\rho 1}$ have weights $\theta^b(r_j)2^t$ on digit t of $e_{j,b}$; $\mathbf{t}_{\rho 2}$ weights 2^t on digits of σ_0 and $\theta^b(r'_j)2^t$ on digits of $e'_{j,b}$; $\mathbf{t}_T$ weights $2^k 2^t$ on digits of σ_0 and $-\eta_X r_i(1-r_i)2^t$ on digits of α_i^X. The honest Prover satisfies

$$\langle \mathbf{t}_\iota, \mathbf{u}_M\rangle = a_\iota \ (\iota \in \mathcal{I} = \{\rho 1, \rho 2, T\}), \qquad a-\mathrm{pub}(\mathbf{r}_y) = \textstyle\sum_i \mathrm{eq}(\mathbf{r}_h, i) g_i, \qquad g_i = \langle \mathbf{b}, \mathbf{U}_i\rangle, \tag{1}$$

with $\mathbf{b} = \big(\mathrm{eq}(\mathbf{r}_l, j)_{j<\gamma} \| (\kappa(\mathbf{r}_l)2^t)_{t<\delta}\big)$. ((T): $\hat{S}' - \sum \eta_X V_X = 2^k\sigma_0 - \sum_X \eta_X \sum_i \alpha_i^X r_i(1-r_i)$.)

Step 4 (mask opening, Σ with aborts + LaBRADOR). $\mathcal{P}$ samples $\mathbf{y} \leftarrow [-B_y^M, B_y^M]^{N_M d}$, $\mathbf{y}_r \leftarrow [-B_y^M, B_y^M]^{l_r^M d}$ and sends $\mathbf{w} = \mathbf{A}_1^M\mathbf{y} + \mathbf{A}_2^M\mathbf{y}_r$, $w_\iota = \langle \mathbf{t}_\iota, \mathbf{y}\rangle$ $(\iota \in \mathcal{I})$. $\mathcal{V} \to \mathcal{P}$: $c \leftarrow \mathcal{C}_R$. $\mathcal{P}$ forms $\mathbf{s} = \mathbf{y} + c\mathbf{u}_M$, $\mathbf{s}_r = \mathbf{y}_r + c\mathbf{r}_M$, restarting if $\|(\mathbf{s}, \mathbf{s}_r)\|_\infty > B_s^M$. It does not send $(\mathbf{s}, \mathbf{s}_r)$ but a LaBRADOR proof of knowledge of them with: $\mathbf{A}_1^M\mathbf{s} + \mathbf{A}_2^M\mathbf{s}_r = \mathbf{w} + c\mathbf{C}_M$; $\langle \mathbf{t}_\iota, \mathbf{s}\rangle = w_\iota + ca_\iota$ $(\iota \in \mathcal{I})$; $\|\mathbf{s}\|^2 + \|\mathbf{s}_r\|^2 \le (\beta_L^M)^2$.

Step 5 (row opening; ν independent repetitions). Public: $\mathbf{b}$, $\mathbf{T}, \mathbf{T}_0, a$. In repetition r: $\mathcal{P}$ samples $\mathbf{Y}, \mathbf{R}_Y \leftarrow [-B_y, B_y]^D$ $(D = (\gamma' + l_r)d)$, computes $\mathbf{C}_Y = \mathbf{A}_1\mathbf{Y} + \mathbf{A}_2\mathbf{R}_Y$, $g_Y = \langle \mathbf{b}, \mathbf{Y}\rangle$, and sends a nonce $\varpi \in \{0,1\}^{2\lambda}$ and $\mathbf{T}_1 = \mathbf{B}_1(\mathbf{G}^{-1}(\mathbf{C}_Y)\|\mathbf{G}^{-1}(g_Y))$. $\mathcal{V} \to \mathcal{P}$: $c_1, \ldots, c_{h+1} \leftarrow \mathcal{S}$. Define $\mathbf{C}_j^{(0)} = \mathbf{C}_j$, $g_j^{(0)} = g_j$, and for $t = 1..h$, $\rho_t = 2^{h-t}$: $\mathbf{C}_j^{(t)} = \mathbf{C}_j^{(t-1)} + c_t\mathbf{C}_{j+\rho_t}^{(t-1)}$, $g_j^{(t)} = g_j^{(t-1)} + c_t g_{j+\rho_t}^{(t-1)}$; $\mathbf{U}^*, \mathbf{R}^*$ analogously from $\mathbf{U}_i, \mathbf{R}_i$. Put $\lambda_i = \prod_t c_t^{i_t}$, $c = c_{h+1}$, and $(\mathbf{V}, \mathbf{R}_V) = (\mathbf{U}^* + c\mathbf{Y}, \mathbf{R}^* + c\mathbf{R}_Y)$; if $\|(\mathbf{V}, \mathbf{R}_V)\|_\infty > B_s = B_y - 2^h B_w$ the repetition restarts. $\mathcal{P}$ sends a LaBRADOR proof of knowledge of $(\mathbf{V}, \mathbf{R}_V, \hat{\mathbf{C}}, \hat{\mathbf{g}}, \hat{\mathbf{C}}_Y, \hat{g}_Y)$ of total ℓ_2-norm $\le \beta_L$ with

(a) $\mathbf{B}\hat{\mathbf{C}} = \mathbf{T}$, (b$_0$) $\mathbf{B}_0\hat{\mathbf{g}} = \mathbf{T}_0$, (b$_1$) $\mathbf{B}_1(\hat{\mathbf{C}}_Y\|\hat{g}_Y) = \mathbf{T}_1$,
(c) $\mathbf{A}_1\mathbf{V} + \mathbf{A}_2\mathbf{R}_V = \sum_i \lambda_i \mathbf{G}\hat{\mathbf{C}}_i + c\,\mathbf{G}\hat{\mathbf{C}}_Y$, (d) $\langle \mathbf{b}, \mathbf{V}\rangle = \sum_i \lambda_i \mathbf{G}\hat{g}_i + c\,\mathbf{G}\hat{g}_Y$,
(e) $a - \mathrm{pub}(\mathbf{r}_y) = \sum_i \mathrm{eq}(\mathbf{r}_h, i)\,\mathbf{G}\hat{g}_i$.

These are dot-product constraints with public ring coefficients (λ_i are monomials), i.e. LaBRADOR's principal relation with all quadratic coefficients zero.

Decision. $\mathcal{V}$ accepts iff all round checks of Steps 2–3 pass, and the LaBRADOR proofs of Step 4 and of all ν repetitions verify. $\mathcal{V}$ computes no fold of commitments; no $\mathbf{C}_i$, g_i or $\mathbf{V}$ is sent.

Non-interactive version. Every challenge is the random oracle H applied to the statement and the transcript so far (including ϖ, so that programmed points have min-entropy $\ge 2\lambda$). Aborting steps are repeated with fresh randomness until the norm test passes.

Proof object. $\pi = \big(\mathbf{T}, \mathbf{C}_M, (\tilde{g}_j), V_\bullet, \hat{S}', (\tilde{g}'_j), a, \mathbf{T}_0,\ \mathbf{w}, w_\iota, \pi_{\mathrm{LB}}^M,\ (\varpi_r, \mathbf{T}_{1,r}, \pi_{\mathrm{LB},r})_{r\le\nu}\big)$.

Remark (design choices). (i) Embedding entries as constants makes all claims R_q-linear without side coefficients. (ii) $\hat{S}'$ is a free uniform message because σ_0 is uniform; (T) ties it to the V_X. (iii) The mask commitment is separate so that the evaluation claim touches only the rows and the keys of the masks have polylogarithmic size. (iv) Repetitions are unbatched so that the repetition theorem [25] applies literally.

4 Security Analysis

Throughout, $q \equiv 5 \pmod 8$, so nonzero $y \in R_q$ with $\|y\|_\infty < \sqrt{q/2}$ are invertible [18]. Let $\varepsilon_{\mathrm{sc}} = (9k+1)/q$, $\varsigma = \sqrt{128/30}$.

4.1 Lemmas

Lemma 3 (Mask consistency). *Let* $\Theta_j = \sum_{b\le 5} e_{j,b}\theta^b$ *and* $\rho_j(T) = \sum_{\mathbf{b}\in\{0,1\}^{k-j-1}} \hat{\rho}_1(r_{<j}, T, \mathbf{b})$. *Then* $\rho_j(T) = 2^{k-j-1}\big(\Theta_j(T) + \sum_{i<j}\Theta_i(r_i)\big)$, $\rho_0(0)+\rho_0(1) = 0$ *and* $\rho_j(0) + \rho_j(1) = \rho_{j-1}(r_{j-1})$. *For* $\hat{\rho}_2$ *and* ρ'_j *the same holds with* $\rho'_j(T) = 2^{k-j-1}(\sigma_0 + \Theta'_j(T) + \sum_{i<j}\Theta'_i(r'_i))$, $\Theta'_j = \sum_{b\le 3} e'_{j,b}\theta^b$, *and* $\rho'_0(0)+\rho'_0(1) = 2^k\sigma_0$.

Proof. Terms with index $> j$ are summed over $\{0,1\}$ and vanish since $\theta^b(0) + \theta^b(1) = 0$; terms with index $< j$ are the constants $\Theta_i(r_i)$, counted 2^{k-j-1} times; the term j gives $2^{k-j-1}\Theta_j(T)$. Using $\Theta_j(0)+\Theta_j(1) = 0$: $\rho_j(0)+\rho_j(1) = 2^{k-j}\sum_{i<j}\Theta_i(r_i) = \rho_{j-1}(r_{j-1})$. The second system is identical, with the constant σ_0 summed to $2^{k-j-1}\sigma_0$.

Lemma 4 (Uniform round messages). *Fix the witness, all masks other than* $e_{j,\cdot}$ *(resp.* $e'_{j,\cdot}$*), and all earlier challenges. If* $e_{j,1..5}$ *is uniform,* $\tilde{g}_j$ *is uniform on* $\mathcal{W}_j = \{h \in \mathbb{F}_q[T]_{\le 5} : h(0)+h(1) = S_j\}$*; if* $e'_{j,1..3}$ *is uniform,* $\tilde{g}'_j$ *is uniform on* $\mathcal{W}'_j = \{h \in \mathbb{F}_q[T]_{\le 3} : h(0)+h(1) = S'_j\}$. *The latter holds for every polynomial* $\hat{z}_{\mathrm{zk}}$ *(in particular for every value of* $\mathbf{p}$*).*

Proof. Write $\tilde{g}_j = g_j^P + \rho_j$ with g_j^P the partial sums of $P = \mathrm{eq}(\boldsymbol{\tau},\cdot)(\tilde{A}_{\mathrm{zk}}\tilde{B}_{\mathrm{zk}} - \tilde{C}_{\mathrm{zk}})$. The sumcheck identity $g_j^P(0) + g_j^P(1) = g_{j-1}^P(r_{j-1})$ (and $= \sum P$ for $j = 0$) with Lemma 3 gives $\tilde{g}_j \in \mathcal{W}_j$. The map $\mathbf{e} \mapsto 2^{k-j-1}\sum_b e_b\theta^b$ is injective (q odd, θ^b independent) onto $\mathcal{D}_5$, the direction space of $\mathcal{W}_j$ (dimension 5), and $g_j^P + 2^{k-j-1}\sum_{i<j}\Theta_i(r_i)$ does not depend on $e_{j,\cdot}$. The second system is the same with $G_2 = M\hat{z}_{\mathrm{zk}}$ in place of P and $\mathcal{D}_3$.

Lemma 5 (Hiding). *If* $3^{l_r d} \ge q^{md}2^{2\lambda}$ *then* $(\mathbf{A}_2, \mathbf{A}_2\mathbf{r})$, $\mathbf{r} \leftarrow \{-1,0,1\}^{l_r d}$, *is* $2^{-\lambda}$*-close to* $(\mathbf{A}_2$, *uniform); so* $\mathbf{A}_1\mathbf{u} + \mathbf{A}_2\mathbf{r}$ *is* $2^{-\lambda}$*-close to uniform for every* $\mathbf{u}$. *The same holds for* $\mathbf{r}$ *uniform on* $[-B,B]^{l_r d}$ *with* $2B < \sqrt{q/2}$ *if* $(2B+1)^{l_r d} \ge q^{md}2^{2\lambda}$.

Proof. For $\mathbf{r} \ne \mathbf{r}'$ the difference has a nonzero coordinate of norm $\le 2B < \sqrt{q/2}$, hence invertible, so $\Pr_{\mathbf{A}_2}[\mathbf{A}_2(\mathbf{r}-\mathbf{r}') = 0] = q^{-md}$. Universality and the leftover hash lemma give distance $\le \frac{1}{2}\sqrt{q^{md}/|\mathrm{support}|} \le 2^{-\lambda-1}$.

Lemma 6 (Aborting response). *Let* $\mathbf{x} \in \mathbb{Z}^D$ *with* $\|\mathbf{x}\|_\infty \le X$, $B_y \ge DX$, $B_s = B_y - X$, $\mathbf{y} \leftarrow [-B_y, B_y]^D$. *Then* $\Pr[\|\mathbf{y}+\mathbf{x}\|_\infty \le B_s] = \left(\frac{2B_s+1}{2B_y+1}\right)^D \ge \frac{1}{4}$, *independently of* $\mathbf{x}$, *and conditioned on this event* $\mathbf{y}+\mathbf{x}$ *is uniform on* $[-B_s, B_s]^D$.

Proof. $\mathbf{y}+\mathbf{x}$ is uniform on a translate of the box by $\mathbf{x}$, which contains $[-B_s, B_s]^D$. The probability is $(1 - \frac{2X}{2B_y+1})^D \ge (1 - \frac{1}{D})^D \ge \frac{1}{4}$ for $D \ge 2$.

For the mask opening $X = \omega B_w$ ($\|c\mathbf{u}_M\|_\infty \le \|c\|_1$, $\omega = 2d$); for the row opening $X = 2^h B_w$ (Step 5 of Theorem 1).

Lemma 7 (Monomials). *(i) For* $s \in \mathcal{S}$: $\|sf\|_\infty = \|f\|_\infty$, $\|sf\|_1 = \|f\|_1$, $s^{-1} \in \mathcal{S}$, *and* $s\mathbf{Y}$ *is uniform on* $[-B, B]^D$ *if* $\mathbf{Y}$ *is. (ii) For distinct* $s, s' \in \mathcal{S}$ *there is* $\bar{g} \in R_q$ *with* $(s - s')\bar{g} = 2$ *and* $\|\bar{g}\|_1 \le d$.

Proof. (i) Multiplication by X^a is a signed permutation of coefficients. (ii) Let $s = X^{a'}$, $s' = X^{a''}$, $a' > a''$ (else swap), $u = a' - a'' = 2^e o$, o odd, $d = 2^\ell$, $e \le \ell$. For $N = 2^{\ell-e}$: $uN = od \equiv d \pmod{2d}$, so $X^{uN} = -1$ and $(X^u - 1)\sum_{i<N} X^{ui} = -2$. Take $\bar{g} = -X^{-a''}\sum_{i<N} X^{ui}$; $\|\bar{g}\|_1 \le N \le d$.

Lemma 8 (Row-masked evaluation). *(i)* $\hat{z}_{\mathrm{zk}}(\mathbf{r}_y) = \mathrm{pub}(\mathbf{r}_y) + \sum_i \mathrm{eq}(\mathbf{r}_h, i) g_i$ *and* $g_i = \langle \mathbf{b}, \mathbf{U}_i \rangle$. *(ii) If* $\kappa(\mathbf{r}_l) \ne 0$ *and* $\mathbf{p}$ *is uniform,* $(g_i)_{i<\rho}$ *is uniform in* $\mathbb{F}_q^\rho$ *and independent of all other transcript elements, for every witness; for digit-defined* p_i *the distance is* $\le \rho 2^{-\lambda}$. *(iii) For arbitrary* $\mathbf{U}_i^* \in \mathbb{F}_q^{\gamma'}$ *with witness parts* W_i^* *and* $p_i^* = \sum_t 2^t U_{i,\gamma+t}^*$, *if* $a - \mathrm{pub}(\mathbf{r}_y) = \sum_i \mathrm{eq}(\mathbf{r}_h, i)\langle \mathbf{b}, \mathbf{U}_i^* \rangle$, *then* $a = \hat{z}_{\mathrm{zk}}^*(\mathbf{r}_y)$ *for* $\hat{z}_{\mathrm{zk}}^* = \hat{z}^* + \kappa\hat{p}^*$, *where* $\hat{z}^*$ *is the extension of the vector* $\mathbf{z}^*$ *built from* (z_P, W^*); *and* $\hat{z}_{\mathrm{zk}}^* = \hat{z}^*$ *on the cube.*

Proof. (i) $\hat{z}(\mathbf{y}) = \sum_i \mathrm{eq}(\mathbf{y}_h, i)\sum_j \mathrm{eq}(\mathbf{y}_l, j) z_{(i,j)}$ and $\hat{p}$ is multilinear; separate the public coordinates. (ii) For fixed $\mathbf{r}_l$, $p_i \mapsto g_i$ is a bijection of $\mathbb{F}_q$ (translation by a constant times $\kappa(\mathbf{r}_l) \ne 0$), the p_i are independent, and $\mathbf{p}$ otherwise enters the transcript through the commitments (Lemma 5) and the second-sumcheck messages, which are uniform for every $\mathbf{p}$ (Lemma 4). $\Pr[\kappa(\mathbf{r}_l) = 0] \le 2/q$. (iii) Same algebra as (i); $\kappa = 0$ on the cube.

Lemma 9 (Sumcheck lemma). *Fix keys and any* $\mathbf{u}_M^* \in \mathbb{F}_q^{N_M}$ *(defining mask scalars* $\alpha^*, e^*, \sigma^*, e'^*$ *by recomposition of digits) and* $\mathbf{U}_i^* \in \mathbb{F}_q^{\gamma'}$ *(defining* $\mathbf{z}^*$ *and* $\mathbf{p}^*$ *as in Lemma 8(iii)). Suppose* $\mathbf{z}^*$ *violates the R1CS. For any adaptive Prover, the probability over* $\boldsymbol{\tau}, \mathbf{r}_x, \boldsymbol{\eta}, \mathbf{r}_y$ *that all round checks pass and*

$$\langle \mathbf{t}_\iota, \mathbf{u}_M^* \rangle = a_\iota \ (\iota \in \mathcal{I}), \qquad a - \mathrm{pub}(\mathbf{r}_y) = \textstyle\sum_i \mathrm{eq}(\mathbf{r}_h, i)\langle \mathbf{b}, \mathbf{U}_i^* \rangle \tag{2}$$

is at most $\varepsilon_{\mathrm{sc}} = (9k+1)/q$.

Proof. Let $\tilde{X}^* = \widetilde{X\mathbf{z}^*}$, $m_X^* = \sum_i \alpha_i^{X*} r_i(1-r_i)$, $\tilde{V}_X^* = \tilde{X}^*(\mathbf{r}_x) + m_X^*$, and G_1^*, G_2^* built from starred objects. (a) $R^* = A\mathbf{z}^* \circ B\mathbf{z}^* - C\mathbf{z}^* \ne 0$, so $\hat{R}^*(\boldsymbol{\tau}) \ne 0$ except with probability k/q; then $\sum_{\mathbf{b}} G_1^*(\mathbf{b}) = \hat{R}^*(\boldsymbol{\tau}) \ne 0 = S_0$. (b) Let $T_j = \sum_{\mathbf{b}} G_1^*(r_{<j}, \mathbf{b})$ and g_j^* the true round polynomial (degree ≤ 5). Invariant: $S_j \ne$

T_j. If $\tilde{g}_j = g_j^*$ the round check fails; otherwise $\tilde{g}_j(r_j) = g_j^*(r_j)$ with probability $\leq 5/q$, else $S_{j+1} \neq T_{j+1}$. So except with probability $5k/q$, $S_k \neq G_1^*(\mathbf{r}_x) = \mathrm{eq}(\boldsymbol{\tau}, \mathbf{r}_x)(\tilde{V}_A^*\tilde{V}_B^* - \tilde{V}_C^*) + \hat{\rho}_1^*(\mathbf{r}_x)$. The Verifier has $S_k = \mathrm{eq}(\ldots)(V_A V_B - V_C) + a_{\rho 1}$ and claim (ρ1) gives $a_{\rho 1} = \hat{\rho}_1^*(\mathbf{r}_x)$; hence $\Delta = \mathbf{V} - \tilde{\mathbf{V}}^* \neq 0$. (c) $\boldsymbol{\eta} \cdot \Delta \neq 0$ except with probability $1/q$. The true value of the second sum is $\hat{S}^* = \sum_X \eta_X \tilde{X}^*(\mathbf{r}_x) + 2^k\sigma^*$. Claim (T) gives $\hat{S}' = \boldsymbol{\eta} \cdot \mathbf{V} + 2^k\sigma^* - \boldsymbol{\eta} \cdot m^* = \hat{S}^* + \boldsymbol{\eta} \cdot \Delta \neq \hat{S}^*$. (d) As in (b) for G_2^* (degree ≤ 3): except with probability $3k/q$, $S'_k \neq G_2^*(\mathbf{r}_y) = M(\mathbf{r}_y)\hat{z}_{\mathrm{zk}}^*(\mathbf{r}_y) + \hat{\rho}_2^*(\mathbf{r}_y)$. By Lemma 8(iii) the row claim gives $a = \hat{z}_{\mathrm{zk}}^*(\mathbf{r}_y)$, and ($\rho$2) gives $a_{\rho 2} = \hat{\rho}_2^*(\mathbf{r}_y)$, so $S'_k = M(\mathbf{r}_y)a + a_{\rho 2} = G_2^*(\mathbf{r}_y)$: contradiction. Total $(k + 5k + 1 + 3k)/q$.

4.2 Completeness

Theorem 1 (Completeness). *An honest Prover with $\|\mathbf{w}\|_\infty \leq B_w$ is accepted, provided each LaBRADOR instance is complete. The expected number of restarts of each aborting step is ≤ 4.*

Proof. Sumchecks. $R = 0$, so $\sum P = 0 = S_0$ and $\sum \hat{\rho}_1 = 0$; by Lemma 4 (membership) all round checks pass and $a_{\rho 1} = \hat{\rho}_1(\mathbf{r}_x)$. $\hat{S}' = \sum G_2$ by definition, Lemma 3 gives the round identities, and $a_{\rho 2} = \hat{\rho}_2(\mathbf{r}_y)$ since $a = \hat{z}_{\mathrm{zk}}(\mathbf{r}_y)$. *Claims.* ($\rho$1), ($\rho$2) follow from the digit recomposition; (T): $\hat{S}' - \boldsymbol{\eta}\cdot\mathbf{V} = 2^k\sigma_0 - \sum_X \eta_X m_X$; the row claims are Lemma 8(i). *Step 4.* By Lemma 6 with $X = \omega B_w$ a response is accepted with probability $\geq 1/4$; linearity gives the verification equations, and $\|\mathbf{s}\|^2 + \|\mathbf{s}_r\|^2 \leq DB_s^2 = (\beta_L^M)^2$. *Step 5.* $\|\mathbf{U}^*\|_\infty \leq 2^h B_w$, $\|\mathbf{R}^*\|_\infty \leq 2^h$ (each fold step adds a monomial multiple of a row, Lemma 7(i)); $c\mathbf{Y}$ is uniform on the box; Lemma 6 gives acceptance $\geq 1/4$ per repetition. Equations (a),(b) hold by definition, (c),(d) by linearity with $\mathbf{G}\hat{\mathbf{C}}_i = \mathbf{C}_i$, $\mathbf{G}\hat{g}_i = g_i$, (e) by Lemma 8(i). The witness norm satisfies $\beta_L^2 \leq DB_s^2 + N_{\mathrm{dig}}2^{2(b-1)}$ with $N_{\mathrm{dig}} = (\rho m_W \tau + \rho\tau + m_W\tau + \tau)d$.

4.3 Knowledge soundness

Lemma 10 (LaBRADOR, Thm. 5.1 of [21], as used). *For the principal relation with norm bound $\beta \leq \sqrt{30/128}\, q/125$, LaBRADOR is a proof of knowledge with knowledge error $\varepsilon_{\mathrm{LB}}$ (2^{-125} in [21]): from a prover accepted with probability $> \varepsilon_{\mathrm{LB}}$ an expected-polynomial-time extractor outputs a witness satisfying every constraint* exactly *with norm $\leq \varsigma\beta$.*

We cite this result and do not reprove it.

Mask opening.

Lemma 11 (Mask extraction). *Fix the prefix π_1 preceding c in Step 4 and let p be the probability (over c and the LaBRADOR round) that Step 4 is accepted. If $p > \eta_1 := \varepsilon_{\mathrm{LB}} + 2/|\mathcal{C}_R|$, an expected-polynomial-time extractor outputs two relaxed openings $(\bar{\mathbf{s}}, \bar{\mathbf{s}}_r, \bar{c})$, $(\bar{\mathbf{s}}', \bar{\mathbf{s}}'_r, \bar{c}')$ of $\mathbf{C}_M$ with ℓ_∞-norm $\leq 2B'$, $B' = \varsigma\beta_L^M$,*

such that $\sum_s t_{\iota,s}\bar{s}_s = \bar{c}a_\iota$. *Moreover, under M-SIS for* $[\mathbf{A}_1^M|\mathbf{A}_2^M]$ *with bound* $8\omega B'$, *the vector* $\mathbf{u}_M^* \coloneqq \mathrm{ct}(\bar{\mathbf{s}}/\bar{c})$ *depends only on* $\mathbf{C}_M$, *and satisfies* $\langle \mathbf{t}_\iota, \mathbf{u}_M^* \rangle = a_\iota$ *for* $\iota \in \mathcal{I}$.

Proof. Let p_c be the acceptance probability for challenge c; $\mathbb{E}_c p_c = p$, so $|\{c : p_c > \varepsilon_{\mathrm{LB}}\}| \geq (p - \varepsilon_{\mathrm{LB}})|\mathcal{C}_R| \geq 2$. For two such $c \neq c'$, Lemma 10 yields exact $(\mathbf{s}, \mathbf{s}_r)$, $(\mathbf{s}', \mathbf{s}_r')$ with norm $\leq B'$. Put $\bar{c} = c - c'$ (nonzero, $\|\bar{c}\|_\infty \leq 4$, invertible, $\|\bar{c}\|_1 \leq 2\omega$): subtracting, $\mathbf{A}_1^M\bar{\mathbf{s}} + \mathbf{A}_2^M\bar{\mathbf{s}}_r = \bar{c}\mathbf{C}_M$ and $\sum t_{\iota,s}\bar{s}_s = \bar{c}a_\iota$. For two relaxed openings of the same $\mathbf{C}_M$, $\mathbf{x} = (\bar{c}^2\bar{\mathbf{s}}^1 - \bar{c}^1\bar{\mathbf{s}}^2,\ \bar{c}^2\bar{\mathbf{s}}_r^1 - \bar{c}^1\bar{\mathbf{s}}_r^2)$ is in the kernel with $\|\mathbf{x}\|_\infty \leq 2 \cdot 2\omega \cdot 2B' = 8\omega B'$; so $\mathbf{x} = 0$ except with M-SIS advantage, i.e. $\bar{\mathbf{s}}/\bar{c}$ is unique. Dividing the claim by $\bar{c}$ gives $\sum t_{\iota,s}\tilde{u}_s = a_\iota$ in R_q for $\tilde{\mathbf{u}} = \bar{\mathbf{s}}/\bar{c}$; $\mathbf{t}_\iota$ and a_ι are constants, so taking constant coefficients gives the claim for $\mathbf{u}_M^*$.

Row opening.

Lemma 12 (Fold extraction with norm recursion). *Let* $X_0 \geq 1$, $X_t = (2d+2)^t X_0$. *Fix commitments* $(\mathbf{C}_i, g_i)_{i<2^h}$ *(*$g_i \in R_q$*),* $\mathbf{C}_Y, g_Y$ *and* $\mathbf{b}$. *Suppose for a full binary tree of challenge tuples* $(c_1, .., c_{h+1})$ *(siblings at depth* t *use distinct* c_t*) there is, at every leaf,* $(\mathbf{V}, \mathbf{R}_V)$ *with* $\|\cdot\|_\infty \leq X_0$, $\mathbf{A}_1\mathbf{V} + \mathbf{A}_2\mathbf{R}_V = \mathbf{C}_{\mathrm{fin}}$, $\langle \mathbf{b}, \mathbf{V} \rangle = g_{\mathrm{fin}}$ *(the folds of the fixed data under the leaf's challenges). Then for every* i *there are* $(\mathbf{V}_i, \mathbf{R}_i^V)$ *with* $\mathbf{A}_1\mathbf{V}_i + \mathbf{A}_2\mathbf{R}_i^V = 2^{h+1}\mathbf{C}_i$, $\langle \mathbf{b}, \mathbf{V}_i \rangle = 2^{h+1}g_i$, *norms* $\leq X_{h+1}$.

Proof. Call $(\mathbf{V}_j, \mathbf{R}_j)_j$ a (μ, X)-opening of rows $(\mathbf{C}_j, g_j)_j$ if $\mathbf{A}\cdot = \mu\mathbf{C}_j$, $\langle \mathbf{b}, \mathbf{V}_j \rangle = \mu g_j$ and norms $\leq X$. We show by induction from the leaves that a node at depth t carries a $(2^{h+1-t}, X_{h+1-t})$-opening of all rows of its level, *with multiplier depending only on depth*. Leaves ($t = h+1$): $(1, X_0)$-openings of the single row $\mathbf{C}_0^{(h)} + c\mathbf{C}_Y$. A node whose children use $s \neq s'$ at the next challenge has two (μ, X)-openings $(\mathbf{V}_j^s), (\mathbf{V}_j^{s'})$ of the rows $\mathbf{C}_j^L + s\mathbf{C}_j^R$, $\mathbf{C}_j^L + s'\mathbf{C}_j^R$ with the *same* μ. With $\bar{g}$ from Lemma 7(ii) put $\mathbf{V}_j^R = \bar{g}(\mathbf{V}_j^s - \mathbf{V}_j^{s'})$, $\mathbf{V}_j^L = 2\mathbf{V}_j^s - s\mathbf{V}_j^R$ (same for $\mathbf{R}$). Then $\mathbf{A}\mathbf{V}_j^R = \bar{g}\mu(s-s')\mathbf{C}_j^R = 2\mu\mathbf{C}_j^R$, $\mathbf{A}\mathbf{V}_j^L = 2\mu(\mathbf{C}_j^L + s\mathbf{C}_j^R) - 2\mu s\mathbf{C}_j^R = 2\mu\mathbf{C}_j^L$; $\langle \mathbf{b}, \cdot \rangle$ is R_q-linear and transforms identically. Norms: $\|\mathbf{V}_j^R\|_\infty \leq d \cdot 2X$, $\|\mathbf{V}_j^L\|_\infty \leq 2X + 2dX$. This is a $(2\mu, (2d+2)X)$-opening of the rows of the parent level. At depth h the two rows are $\mathbf{C}_0^{(h)}$ and $\mathbf{C}_Y$; we keep the first. After $h+1$ levels we obtain a $(2^{h+1}, X_{h+1})$-opening of $(\mathbf{C}_i, g_i)_i$.

No relaxation factor from a Σ-protocol appears at the leaves, because LaBRADOR's extractor outputs an exact witness; this is why multipliers add rather than multiply.

Lemma 13 (Binding). *Two* (μ, X)*-openings of the same commitment with the same integer multiplier* μ *coincide unless a nonzero M-SIS solution of norm* $\leq 2X$ *is found. In particular two exact openings (*$\mu = 1$, *norm* $\leq X_0$*) of* $\mathbf{T}$, $\mathbf{T}_0$ *or* $\mathbf{T}_1$ *coincide unless M-SIS*$_{2X_0}$ *is broken for* $\mathbf{B}, \mathbf{B}_0, \mathbf{B}_1$*; two* $(2^{h+1}, X_{h+1})$*-openings of* $\mathbf{C}_i$ *coincide unless M-SIS*$_{2X_{h+1}}$ *is broken for* $[\mathbf{A}_1|\mathbf{A}_2]$.

Proof. The difference lies in the kernel and has norm $\leq 2X$.

Theorem 2 (Special soundness of one repetition). *Fix the statement* $(\mathbf{T}, \mathbf{T}_0, a,$ *sumcheck transcript*$)$ *and* $\mathbf{T}_1, \varpi$*. Suppose a full binary tree of depth* $h+1$ *in the challenges of one repetition has accepting leaves with exact LaBRADOR-extracted witnesses* ($\ell_\infty \le X_0 = \varsigma\beta_L$)*. Then, except with the advantage of an M-SIS adversary against the keys, there are* $\mathbf{U}_i^* \in \mathbb{F}_q^{\gamma'}$ *(*$i < \rho$*) with (1)* $\mathbf{U}_i^*$ *is determined by* $\mathbf{T}$ *alone; (2)* $a - \mathrm{pub}(\mathbf{r}_y) = \sum_i \mathrm{eq}(\mathbf{r}_h, i)\langle \mathbf{b}, \mathbf{U}_i^* \rangle$.

Proof. All leaves share $\mathbf{T}, \mathbf{T}_0, \mathbf{T}_1$, so by (a), ($\mathrm{b}_0$), ($\mathrm{b}_1$) and Lemma 13 they carry the same $\hat{\mathbf{C}}, \hat{\mathbf{g}}, \hat{\mathbf{C}}_Y, \hat{g}_Y$. Set $\mathbf{C}_i = \mathbf{G}\hat{\mathbf{C}}_i$, $g_i = \mathbf{G}\hat{g}_i \in R_q$, $\mathbf{C}_Y = \mathbf{G}\hat{\mathbf{C}}_Y$, $g_Y = \mathbf{G}\hat{g}_Y$: fixed for the whole tree. At each leaf (c),(d) are the hypotheses of Lemma 12, which yields $(2^{h+1}, X_{h+1})$-openings $(\mathbf{V}_i, \mathbf{R}_i^V)$ of $(\mathbf{C}_i, g_i)$. As $\mathbf{C}_i$ is determined by $\mathbf{T}$, Lemma 13 makes $\mathbf{V}_i$ determined by $\mathbf{T}$. Put $\mathbf{U}_i^* = \mathrm{ct}(\mathbf{V}_i)/2^{h+1}$ (2 invertible). Taking constant coefficients of $\langle \mathbf{b}, \mathbf{V}_i \rangle = 2^{h+1} g_i$ ($\mathbf{b}$ constant) gives $\langle \mathbf{b}, \mathbf{U}_i^* \rangle = \mathrm{ct}(g_i)$; taking constant coefficients of (e) (constant public coefficients) gives $a - \mathrm{pub} = \sum_i \mathrm{eq}(\mathbf{r}_h, i)\,\mathrm{ct}(g_i)$.

Lemma 14 (Row-layer knowledge soundness). *Fix the prefix* π_2 *preceding the challenges of Step 5 and let* p *be the probability that all* ν *repetitions are accepted. If* $p > \eta_2 := \kappa_{\mathrm{f}}^\nu + \nu\varepsilon_{\mathrm{LB}}$, $\kappa_{\mathrm{f}} = 1 - (1 - \frac{1}{2d})^{h+1} \le \frac{h+1}{2d}$*, an expected-polynomial-time extractor outputs* $(\mathbf{U}_i^*)_i$ *as in Theorem 2.*

Proof. Consider the clear variant of one repetition (final message: the witness itself). By Theorem 2 it is $(2, \ldots, 2)$-special-sound with $h+1$ challenge rounds from $\mathcal{S}$, $|\mathcal{S}| = 2d$, so its ν-fold parallel repetition has knowledge error κ_{f}^ν by [25]. The real protocol replaces the witness of each repetition by a LaBRADOR proof; the LaBRADOR extractor (Lemma 10) applied to each repetition's residual strategy converts the prover into one for the clear variant, losing $\varepsilon_{\mathrm{LB}}$ per repetition by a union bound. The expected-time bookkeeping of composing the two extractors is standard and not quantified here; we use [25] through its abstract-level statement and the reader should verify its exact hypotheses.

Theorem 3 (Knowledge soundness, interactive). *Let* $\mathcal{P}^*$ *be accepted with probability* ϵ*. Under M-SIS for the five keys with the bounds above (total advantage* Adv*), if* $\epsilon > \varepsilon_{\mathrm{sc}} + \eta_1 + \eta_2 + \mathrm{Adv}$ *there is an expected-polynomial-time extractor that outputs* $\mathbf{w}^* \in \mathbb{F}_q^{n_w}$ *such that* $(z_P, \mathbf{w}^*)$ *satisfies the R1CS.*

Proof. Let $\varphi = (\mathbf{T}, \mathbf{C}_M)$ be the first message. Lemma 11 and Lemma 14 define, for every prefix on which their thresholds are exceeded, $\mathbf{u}_M^*$ and $(\mathbf{U}_i^*)_i$; by their uniqueness parts (except for Adv) these depend on φ only, hence are fixed before $\boldsymbol{\tau}$ is drawn. Let $\mathbf{z}^*$ be the R1CS vector built from $(z_P, \mathbf{U}^*)$. Order the protocol so that Step 4 precedes Step 5 and let π_1, π_2 be the prefixes before the challenge of Step 4 and Step 5. Let $p_1(\pi_1)$, $p_2(\pi_2)$ be the acceptance probabilities of Step 4 given π_1 and of Step 5 given π_2. An accepted run with $p_1(\pi_1) \le \eta_1$ has probability mass at most η_1 (conditioned on π_1, acceptance has probability $\le p_1$); likewise for $p_2(\pi_2) \le \eta_2$. Hence the set $\mathcal{G}$ of accepted runs with $p_1 > \eta_1$ and $p_2 > \eta_2$ has mass $\ge \epsilon - \eta_1 - \eta_2$. On a run in $\mathcal{G}$ the round checks of Steps 2–3 passed, and Lemmas 11, 14 give that (2) holds for $(\mathbf{u}_M^*, \mathbf{U}^*)$ with this run's challenges. If

$\mathbf{z}^*$ violated the R1CS, Lemma 9 would bound the mass of such runs by $\varepsilon_{\mathrm{sc}}$; so $\epsilon \le \eta_1 + \eta_2 + \varepsilon_{\mathrm{sc}} + \mathrm{Adv}$, a contradiction. The extractor samples accepted runs, applies the two extractors at π_1, π_2, and outputs the witness part of $\mathbf{U}^*$.

The extractor obtains an arbitrary vector in $\mathbb{F}_q^{n_w}$, not a short one. Boolean witnesses require added constraints $z_i^2 = z_i$.

4.4 Zero-knowledge

Theorem 4 (Zero-knowledge, programmable random oracle). *There is a simulator that, given only the statement and key, with programmable H, outputs a proof whose joint distribution with H is within statistical distance*

$$(11k + 2 + 2\rho)\, 2^{-\lambda} + 2(2/q)^k + 2/q + (\nu + 1)Q\, 2^{-2\lambda}$$

of the real one, against any distinguisher making Q classical queries.

Proof. **Simulator.** (1) $\mathbf{C}_i \leftarrow R_q^m$, $\mathbf{T} = \mathbf{B}\mathbf{G}^{-1}(\mathbf{C}_\bullet)$, $\mathbf{C}_M \leftarrow R_q^{m_M}$. (2) Sample $\tilde{g}_j \leftarrow \mathcal{W}_j$ uniformly, obtain r_j from H, set S_{j+1}; sample $V_X \leftarrow \mathbb{F}_q$; $a_{\rho 1} = S_k - \mathrm{eq}(V_A V_B - V_C)$. Obtain $\boldsymbol{\eta}$, sample $\hat{S}' \leftarrow \mathbb{F}_q$, $\tilde{g}'_j \leftarrow \mathcal{W}'_j$ (with $S'_0 = \hat{S}'$), obtain $\mathbf{r}_y$, sample $a \leftarrow \mathbb{F}_q$, set $a_{\rho 2} = S'_k - M(\mathbf{r}_y)a$, $a_T = \hat{S}' - \boldsymbol{\eta} \cdot \mathbf{V}$, and g_i uniform with $\sum \mathrm{eq}(\mathbf{r}_h, i) g_i = a - \mathrm{pub}(\mathbf{r}_y)$; $\mathbf{T}_0$ from g. (3) Step 4: $c \leftarrow \mathcal{C}_R$, $(\mathbf{s}, \mathbf{s}_r) \leftarrow [-B_s^M, B_s^M]$, $\mathbf{w} = \mathbf{A}^M \cdot (\mathbf{s}, \mathbf{s}_r) - c\mathbf{C}_M$, $w_\iota = \langle \mathbf{t}_\iota, \mathbf{s} \rangle - c a_\iota$, program H at the point of $(\mathbf{w}, w_\bullet)$ to return c; run the honest LaBRADOR prover. (4) Step 5, for each repetition: pick $c_1, .., c_{h+1} \leftarrow \mathcal{S}$ and $(\mathbf{V}, \mathbf{R}_V) \leftarrow [-B_s, B_s]^D$; set $\mathbf{C}_Y = c^{-1}(\mathbf{A}_1\mathbf{V} + \mathbf{A}_2\mathbf{R}_V - \sum_i \lambda_i \mathbf{C}_i)$, $g_Y = c^{-1}(\langle \mathbf{b}, \mathbf{V} \rangle - \sum_i \lambda_i g_i)$; pick a random nonce ϖ, form $\mathbf{T}_1$, program H on the whole chain $(\mathbf{T}_1, \varpi) \mapsto c_1 \mapsto c_2 \cdots \mapsto c_{h+1}$; run the honest LaBRADOR prover on the witness so defined. A programmed point was queried before with probability $\le Q 2^{-2\lambda}$ per programming (nonce; for Step 4 the min-entropy of $\mathbf{w}$: if an entry of $\mathbf{A}_2^M$ is invertible, $s_{r,1} \mapsto \mathbf{w}$ is injective, so $\mathbf{w}$ has min-entropy $\ge d \log_2(2B_s^M + 1) \ge 2\lambda$).

Hybrids. H_0 real. H_1 *(responses).* Replace Steps 4 and 5 by the simulator's: in the real run each attempt has fresh randomness and a fresh oracle point, so c is uniform and independent of $(\mathbf{y}, \mathbf{y}_r)$; by Lemma 6 the accepted $(\mathbf{s}, \mathbf{s}_r)$ is uniform on the box and independent of $(\mathbf{u}_M, \mathbf{r}_M, c)$, and $\mathbf{w}, w_\iota$ are determined by $(c, \mathbf{s}, \mathbf{s}_r)$ and a_ι. In Step 5, by Lemmas 6, 7(i), the accepted $(\mathbf{V}, \mathbf{R}_V)$ is uniform and independent of the witness, and $\mathbf{C}_Y, g_Y$ are the stated functions of $(\mathbf{V}, \mathbf{R}_V, c, \mathbf{C}_i, g_i)$. H_2 *(LaBRADOR).* In both worlds LaBRADOR runs on identically distributed inputs, so its (absent) zero-knowledge is irrelevant; this uses that after H_1 the secrets $(\mathbf{u}_M, \mathbf{r}_M, \mathbf{R}_i)$ occur only in $\mathbf{C}_M, \mathbf{C}_i$ and in the claim values. H_3 *(commitments).* Replace $\mathbf{C}_M$ and each $\mathbf{C}_i$ by uniform: $(1+\rho)2^{-\lambda}$ (Lemma 5). H_4 *(digits).* Mask scalars and p_i become exactly uniform: $(11k + 1 + \rho)2^{-\lambda}$. H_5 *(sumcheck 1).* Conditioned on all other masks and the witness, $e_{j,\cdot}$ is uniform, so each $\tilde{g}_j$ is uniform on $\mathcal{W}_j$ (Lemma 4); exact. H_6 *(V_X).* If some $r_i \notin \{0, 1\}$ (failure $(2/q)^k$), $V_X = \tilde{X}(\mathbf{r}_x) + \sum \alpha_i^X r_i(1 - r_i)$ is uniform and independent of the view. H_7 *($\hat{S}'$).* $\hat{S}' = \sum \eta_X \tilde{X}(\mathbf{r}_x) + 2^k \sigma_0$ with σ_0 otherwise unused so far: uniform. H_8 *(sumcheck 2).* $\tilde{g}'_j$ uniform on $\mathcal{W}'_j$ for every value of σ_0 and $\mathbf{p}$ (Lemma 4). H_9 *(a, g).* By

Lemma 8(ii) g is uniform and independent of the view, up to $\Pr[\kappa(\mathbf{r}_l) = 0] \le 2/q$; $a = \text{pub} + \sum \text{eq}\, g$ is then uniform, and given a, g is uniform on the hyperplane $\sum \text{eq}(\mathbf{r}_h, i) g_i = a - \text{pub}$ (nonempty since $\sum_i \text{eq}(\mathbf{r}_h, i) = 1$). This is the simulator's step (2). The remaining quantities $a_{\rho 1}, a_{\rho 2}, a_T$ are functions of the view. Summing the terms gives the stated bound.

Remark 1. For quantum adversaries the programming steps would need adaptive reprogramming bounds [23] and the Fiat–Shamir-with-aborts analysis of [19]; we claim zero-knowledge only classically.

4.5 Non-interactive soundness

Theorem 3 concerns the interactive public-coin protocol. We do not prove soundness of the Fiat–Shamir compiled protocol. The sumcheck rounds and the monomial-fold rounds are special-sound (degrees ≤ 5, ≤ 3; 2 per fold round), so existing results for multi-round special-sound protocols [24, 20] apply to those parts; LaBRADOR's own Fiat–Shamir soundness is a standing heuristic also in [21]. Our non-interactive claim is therefore conditional on that heuristic and on the loss terms of [24], which we do not compute.

5 Parameters and Efficiency

Estimating M-SIS. For $R = md$ rows, cols columns and ℓ_∞-bound β we use the SIS instance with ℓ_2-bound $\beta\sqrt{\text{cols}}$ (ignoring ring structure). BKZ with block size b finds vectors of norm $\delta(b)^w q^{R/w}$ in a q-ary lattice of dimension $w \le$ cols, $\delta(b) = \big((\pi b)^{1/b} b/(2\pi e)\big)^{1/(2(b-1))}$, provided $\beta\sqrt{\text{cols}}$ exceeds the Gaussian-heuristic length; we scan w and report the smallest b. Cost is $2^{0.292b}$ (classical) and $2^{0.265b}$ (quantum) in the Core-SVP model. The code reproduces $b = 416$ for the Falcon-512 SIS instance (specification: ≈ 411). We require $b \ge 490$ (> 129 quantum bits). The estimate excludes ring-specific and hybrid attacks.

Parameter set. $n = 2^{20}$ ($k = 20$), $\lambda = 128$, $B_w = 1$, $d = 64$, $q \approx 2^{160}$ prime $\equiv 5 \pmod 8$ (so $\varepsilon_{\text{sc}} \approx 2^{-152}$), $\delta = 288$, $h = 7$ ($\rho = 128$, $\gamma' = 8480$), gadget base 2^{32} ($\tau = 5$), $\nu = \lceil 128/(-\log_2 \kappa_{\text{f}}) \rceil = 32$ with $\kappa_{\text{f}} \le 8/128$. Then $D = (\gamma' + l_r)d \approx 9.3 \times 10^5$ with $l_r = 6060$, $B_y = D2^h \approx 2^{26.8}$, $\beta_L \approx 2^{36.7}$ for the $\mathbf{V}$-part and $X_0 = \varsigma\beta_L$ including the digit part $\approx 2^{42.9}$, $\log_2 \beta_{\text{row}} = \log_2(2(2d+2)^{h+1} X_0) \approx 100$. LaBRADOR's requirement $\beta \le \sqrt{30/128}\, q/125$ holds with a wide margin.

Proof size. Ring elements: $\mathbf{C}_M, \mathbf{w}$ ($2 \cdot 19$), w_ι (3), $\mathbf{T}, \mathbf{T}_0$ ($15+14$), and ν commitments $\mathbf{T}_1$ ($32 \cdot 13$): 486 elements of $d \log_2 q/8 = 1280$ bytes, i.e. ≈ 0.62 MB. Field elements: $5k+3k$ round coefficients (using the sum constraint), $V_A, V_B, V_C, \hat{S}', a$: ≈ 3.3 KB. Nonces: $\nu \cdot 32$ bytes. Plus $\nu + 1 = 33$ LaBRADOR proofs, each polylogarithmic in the witness dimension [21]. We do not estimate their size for our relations; [21] reports tens of kilobytes at comparable instance sizes, which is an order of magnitude only.

Table 2: M-SIS instances for the five keys at the parameter set above (b = estimated BKZ block size).

Key	Role	m	columns	$\log_2 \beta_\infty$	b
$[\mathbf{A}_1 \mid \mathbf{A}_2]$	rows, bound $2X_{h+1}$	60	9.3×10^5	100.0	498
$\mathbf{B}$	$\mathbf{T}$, bound $2X_0$	15	2.5×10^6	43.9	511
$\mathbf{B}_0$	$\mathbf{T}_0$	14	4.1×10^4	43.9	543
$\mathbf{B}_1$	$\mathbf{T}_1$	13	2.0×10^4	43.9	504
$[\mathbf{A}_1^M \mid \mathbf{A}_2^M]$	mask, bound $8\omega\varsigma\beta_L^M$	19	4.2×10^6	51.0	494

Verifier. $O(k)$ for the sumchecks; T_{circ} plus pub (Lemma 2); $O(k\delta)$ for the mask functionals; and the LaBRADOR verifications, linear in their statements. Per repetition the statement contains $\mathbf{A}$ ($m_W \times \gamma'$), the fold coefficients ($m_W \rho\tau$), $\mathbf{B}$ ($m_B \times \rho m_W \tau$), $\mathbf{B}_0, \mathbf{B}_1$, giving

$$\nu\big[m_W(\gamma' + \rho\tau)d + m_B \rho m_W \tau d + m_0 \rho\tau d + m_1(m_W\tau + \tau)d\big] \approx 2.3 \times 10^9$$

coefficient operations at $n = 2^{20}$, dominated by the $\mathbf{A}$ term ($\approx 1.0 \times 10^9$) and the fold terms. This exceeds $O(\text{nnz}) \approx 10^6$ for direct checking. With $\rho \propto \sqrt{n}$ and ignoring slow growth of m, ν, the cost scales as $\sqrt{n}$ and falls below n for $n \gtrsim 2^{42}$. Alternatives in the same family (other d, q, h) give 1.7–2.3×10^9; larger h shifts the cost from $\mathbf{A}$ to $\mathbf{B}$ without reducing it below $\approx 0.9 \times 10^9$.

Prover. Forming $A\mathbf{z}, B\mathbf{z}, C\mathbf{z}$ costs $O(\text{nnz})$ and the sumchecks $O(n)$ field operations (tables folded in place; mask terms in closed form). Committing the rows costs $m_W\, n\, d \approx 4 \times 10^9$ coefficient operations because entries are constants of R_q (the price of avoiding leaked side coefficients). Each repetition runs LaBRADOR on $\approx 5 \times 10^4$ ring elements (3×10^6 coefficients), ν times, plus the mask instance of $\approx 4 \times 10^6$ coefficients. Expected restarts are ≤ 4 per aborting step.

Table 3: Asymptotics, n = circuit size. Proof size counts the $\nu = O(\lambda)$ factor but not poly(λ).

Scheme	Proof size	Verifier	PQ	Transp.
Groth16 [2]	$O(1)$	$O(1)$	×	×
STARKs [3]	$O(\log^2 n)$	$O(\log^2 n)$	✓	✓
Bulletproofs [4]	$O(\log n)$	$O(n)$	×	✓
Ling et al. [1]	$O(\lambda n)$	$O(n)$	✓	✓
LaBRADOR [21] (not ZK)	$O(\log n)$	$O(n)$	✓	✓
This work	polylog(n)	$\tilde{O}(\sqrt{n}) + T_{\text{circ}}$	✓	✓

6 Conclusion

We constructed a transparent post-quantum zero-knowledge argument for R1CS with polylogarithmic proof size and an $\tilde{O}(\sqrt{n})$ verifier up to the circuit-evaluation term, and proved completeness, interactive knowledge soundness and classical zero-knowledge. The sublinear gain is asymptotic: at $n = 2^{20}$ the verifier costs about 2×10^9 coefficient operations and proofs carry $\approx 0.6\,\mathrm{MB}$ of commitments plus 33 LaBRADOR proofs. The cost is driven by the proven norm bounds and by the ν-fold repetition forced by the small monomial challenge set. Open: soundness of the Fiat–Shamir compilation (including LaBRADOR's), a quantum-oracle zero-knowledge bound, an evaluator beyond data-parallel circuits, a quantified composition of the tree and LaBRADOR extractors, and challenge sets larger than $2d$ with short-inverse differences, which would remove ν and reduce the verifier cost by roughly a factor of 30.

References

1. S. Ling, K. Nguyen, D. H. Phan, H. Tang, H. Wang. Zero-knowledge proofs for committed symmetric Boolean functions. *PQCrypto 2021*, LNCS 12841, pp. 339–359.
2. J. Groth. On the size of pairing-based non-interactive arguments. *EUROCRYPT 2016*, LNCS 9666, pp. 305–326.
3. E. Ben-Sasson, I. Bentov, Y. Horesh, M. Riabzev. Scalable zero knowledge with no trusted setup. *CRYPTO 2019*, LNCS 11694, pp. 701–732.
4. B. Bünz, J. Bootle, D. Boneh, A. Poelstra, P. Wuille, G. Maxwell. Bulletproofs: Short proofs for confidential transactions and more. *IEEE S&P 2018*, pp. 315–334.
5. J. Bootle, V. Lyubashevsky. Efficient lattice-based inner-product arguments. *EUROCRYPT 2022*.
6. O. Regev. On lattices, learning with errors, random linear codes, and cryptography. *STOC 2005*, pp. 84–93.
7. D. Micciancio, O. Regev. Worst-case to average-case reductions based on Gaussian measures. *FOCS 2004*, pp. 372–381.
8. J. Don, S. Fehr, C. Majenz, C. Schaffner. Security of the Fiat-Shamir transformation in the quantum random-oracle model. *CRYPTO 2019*, LNCS 11693, pp. 356–383.
9. A. Fiat, A. Shamir. How to prove yourself: Practical solutions to identification and signature problems. *CRYPTO 1986*, LNCS 263, pp. 186–194.
10. S. Goldwasser, S. Micali, C. Rackoff. The knowledge complexity of interactive proof systems. *SIAM J. Comput.*, 18(1):186–208, 1989.
11. A. Gabizon, Z. J. Williamson, O. Ciobanu. PLONK: Permutations over Lagrange-bases for oecumenical noninteractive arguments of knowledge. *IACR ePrint 2019/953*, 2019.
12. P. W. Shor. Polynomial-time algorithms for prime factorization and discrete logarithms on a quantum computer. *SIAM J. Comput.*, 26(5):1484–1509, 1997.
13. A. Kate, G. M. Zaverucha, I. Goldberg. Constant-size commitments to polynomials and their applications. *ASIACRYPT 2010*, LNCS 6477, pp. 177–194.
14. S. Setty. Spartan: Efficient and general-purpose zkSNARKs without trusted setup. *CRYPTO 2020*, LNCS 12172, pp. 704–737.

15. C. Lund, L. Fortnow, H. Karloff, N. Nisan. Algebraic methods for interactive proof systems. *J. ACM*, 39(4):859–868, 1992.
16. S. Ames, C. Hazay, Y. Ishai, M. Venkitasubramaniam. Ligero: Lightweight sublinear arguments without a trusted setup. *ACM CCS 2017*, pp. 2087–2104.
17. C. Baum, I. Damgård, V. Lyubashevsky, S. Oechsner, C. Peikert. More efficient commitments from structured lattice assumptions. *SCN 2018*, LNCS 11035, pp. 368–385.
18. V. Lyubashevsky, G. Seiler. Short, invertible elements in partially splitting cyclotomic rings and applications to lattice-based zero-knowledge proofs. *EUROCRYPT 2018*, LNCS 10820, pp. 204–224.
19. E. Kiltz, V. Lyubashevsky, C. Schaffner. A concrete treatment of Fiat-Shamir signatures in the quantum random-oracle model. *EUROCRYPT 2018*, LNCS 10822, pp. 552–586.
20. J. Don, S. Fehr, C. Majenz, C. Schaffner. Online-extractability in the quantum random-oracle model. *EUROCRYPT 2022*, LNCS 13277, pp. 677–706.
21. W. Beullens, G. Seiler. LaBRADOR: Compact proofs for R1CS from Module-SIS. *CRYPTO 2023*, LNCS 14085, pp. 518–548.
22. N. K. Nguyen, G. Seiler. Greyhound: Fast polynomial commitments from lattices. *CRYPTO 2024*.
23. A. Grilo, A. Hülsing, K. Hövelmanns, C. Majenz. Tight adaptive reprogramming in the QROM. *ASIACRYPT 2021*, LNCS 13090, pp. 637–667.
24. T. Attema, S. Fehr, M. Kloß. Fiat-Shamir transformation of multi-round interactive proofs. *TCC 2022*, LNCS 13747, pp. 113–142.
25. T. Attema, S. Fehr. Parallel repetition of $(k_1, \ldots, k_\mu)$-special-sound multi-round interactive proofs. *CRYPTO 2022*; ePrint 2021/1259.